Telestream

# DIVA

# Operations Guide

**Release: 9.0**

**Revision: 1.0**

**January 9, 2024**

24010091355

# Copyrights and Trademark Notices

# Contents

# Telestream Contact Information

To obtain product information, technical support, or provide comments on this guide, contact us using our web site, email, or phone number as listed below.

| Resource | Contact Information |
|---|---|
| DIVA Technical Support | Web Site:<br>https://www.telestream.net/telestream-support/<br>Depending on the problem severity, Telestream will respond to your request within 24 business hours. For P1, Telestream will respond within 1 hour. Please see the Maintenance & Support Guide for definitions.<br>• Support hours for customers are Monday through Friday, 7 AM to 6 PM, local time.<br>• Support for P1 issues for customers are 24/7. |
| Telestream, LLC | Web Site: www.telestream.net<br>Sales and Marketing Email: info@telestream.net<br>Telestream, LLC<br>848 Gold Flat Road, Suite 1<br>Nevada City, CA USA 95959 |
| International Distributor Support | Web Site: www.telestream.net<br>See the Telestream Web site for your regional authorized Telestream distributor. |
| Telestream Technical Writers | Email: techwriter@telestream.net<br>Share comments about this or other Telestream documents. |

# Preface

This book outlines general operational guidelines for DIVA—often referred to simply as *TSCC*. Included are start-up and shut-down procedures for various software and hardware components of a typical TSCC system, and the control and monitoring aspects of the TSCC platform using the web app.

## Topics

- Audience
- Documentation Accessibility
- Related Documents
- Document Updates

# Audience

This book is intended for operations and administration personnel.

# Documentation Accessibility

For information about Telestream's commitment to accessibility, visit the Telestream Support Portal located at https://www.telestream.net/telestream-support/

# Related Documents

For more information, see the DIVA documentation set located at:

https://www.telestream.net/telestream-support/

For information on Cloud Storage visit these links:

## Metered and non-metered Oracle Cloud Storage

http://docs.oracle.com/en/cloud/get-started/subscriptions-cloud/csgsg/

## Up to date Oracle Cloud information

http://docs.oracle.com/cloud/latest/

## EMC ECS (Elastic Cloud Storage)

https://www.delltechnologies.com/ru-by/learn/data-storage/ecs.htm

## Amazon S3 Cloud Storage

https://aws.amazon.com/s3/

## Scality Zenko Integration

https://www.zenko.io/what-is-zenko/

## Cloudian

https://cloudian.com/

## NetApp StorageGrid

https://www.netapp.com/cloud-services/

## Alibaba OSS

https://www.alibabacloud.com/product/oss

# Document Updates

The following table identifies updates made to this document.

| Date | Update |
|---|---|
| March 2023 | Updated book to change product name from DIVA Core to Content Manager. |
| August 2023 | • Removed information to create new book for the 9.0 release.<br>• Updated book with all new information from the 8.3.1 release.<br>• Removed references to Oracle database; replaced with Postgres.<br>• Change product name from Content Manager to DIVA. |

# Overview

The DIVA architecture enables integration of many different types of servers and technologies—for example, broadcast video servers, storage area networks, and enterprise tape managed storage. This chapter provides an overview of these features, with details.

## Topics

- Overview
- New and Enhanced Features and Functionality

# Overview

The installation of DIVA varies from site to site, so the exact configuration of your specific DIVA platform is not covered in this guide. Refer to the System Administrator, the Telestream Installation and Delivery Team, or Technical Support for details on your specific DIVA System installation and configuration.

# New and Enhanced Features and Functionality

Refer to the DIVA Release Notes located at:

https://www.telestream.net/telestream-support/

## DIVA Command, Configuration Utility, and Control Panel

The DIVA Command, DIVA Configuration Utility, and the DIVA Control Panel have been deprecated and are replaced with the DIVA web app.

# Concepts

To operate and utilize DIVA successfully, it is helpful to understand the different components and concepts of DIVA. Also, refer to the DIVA Architecture, Concepts and Glossary Guide on the DIVA Technical Support site.

## Topics

- Unmanaged Storage Repositories
- Cloud Replicated Bucket Scanning
- Arrays, Disks, and Cache
- Tape Groups and Sets
- Media Storage Formats
- Objects
- Jobs
- Job Types
- Metasources
- Symbolic Links
- Storage Policy Management
- Checksum Support and Content Verification
- Quality of Service

# Unmanaged Storage Repositories

A Source Unmanaged Storage Repository is defined as any connected system that has content intended to be transferred to DIVA. A Destination Unmanaged Storage Repository is defined as any connected system that requires content to be transferred to it from DIVA. Examples of both are Broadcast Video Servers, FTP Servers, or Disk Storage.

UNC paths are supported for SMB Servers and managed disks if the UNC path is directly mounted on the Windows Actors.

The Source and Destination Unmanaged Storage Repositories that are used in DIVA jobs are predefined in the DIVA configuration and are accessible through the web app’s Resource Management menu item In DIVA’s server configuration, each server type or disk file system is given a unique name and are configured as follows:

- **Source Only**

  DIVA only archives files from the server or disk file system.

- **Destination Only**

  DIVA only restores files to the server or disk file system.

- **Source and Destination**

  DIVA archives and restores files to and from the server or disk file system.

Although a detailed explanation of the configuration of a Source or Destination Server is beyond the scope of this guide, a brief overview of the configuration is included because it can influence how jobs are issued to them, and influence how two or more simultaneous jobs to them are managed in the Jobs queue.

Generally, each Source and Destination Unmanaged Storage Repository uses the following parameters. These are common to all jobs that involve the target Server:

- The Source Type is the protocol or access method used when interacting with the target device.
- The maximum number of read and write transfer sessions and the total maximum number of read/write sessions combined. This identifies the limits on the number of simultaneous jobs that DIVA executes concurrently on the target device, or prioritizing write (Restore) operations over read (Archive) operations.
- Define the maximum bandwidth allowable to DIVA for transfers to or from the device. This may be used to throttle data transfers where the target device is shared with other Networks or third party applications.
- The Default Quality of Service (QOS). This is the QOS used when Default is specified in a job's Quality of Service field.
- Define Connect Options that must be provided (or that can also be optionally specified) for the specific protocol or access method of the target device. Examples of Connect Options are recursive subfolders, user names or passwords, or other options specific to the selected source type. DIVA ignores this parameter if no options are specified.

- The Root Path to the files to be archived on the source, or restored to on a destination. This is always specified as an absolute directory path on the target device. For example c:\Exported\MPEG2 for Windows based file systems, or /Movies/MPEG2 on Linux based file systems. The Root Path configuration also depends on the source type, and can be left blank in some cases (and is ignored by DIVA).

  For Local or Disk source types, the Root Path specifies the mount point of the device in the local file system.

If the Connect Options and Root Path parameters have been defined for a Server configuration, they may not be appropriate for every job submitted. DIVA allows these parameters to be specified in a DIVA job to that source or destination (at the job level). Whether a job can override these Server attributes depends on the source type. See the DIVA Source and Destination Servers table in the DIVA Installation and Configuration Guide (https://www.telestream.net/telestream-support/) for a comprehensive list of these options, paths, and how they interact with those specified at the job level.

The Files Path Root specified in a job can either be appended to the Root Path specified in the Server configuration, or override the Root Path entirely if it is specified as an absolute path.

# Alibaba OSS

DIVA includes support for Alibaba OSS integration as both Storage Accounts and Servers. Storage accounts allow users to configure programmatic access to a user's account. The configuration data in a DIVA storage account is used by the DIVA Actors to query storage and transfer content. Storage Class supported = STANDARD. ARCHIVE bucket support is planned for a future DIVA version.

Refer to the DIVA Installation and Configuration Guide for more details on these storage and Servers.

# Alto Disk Archive Integration

Alto Disk Archive is a type of library of disks. Instead of mounting/dismounting tape, this library is designed to mount/dismount the file system of unpluggable disks and map them to an SMB share.

DIVA supports Alto like a tape library. A disk of the library is seen as a tape in DIVA with its UUID as the label.

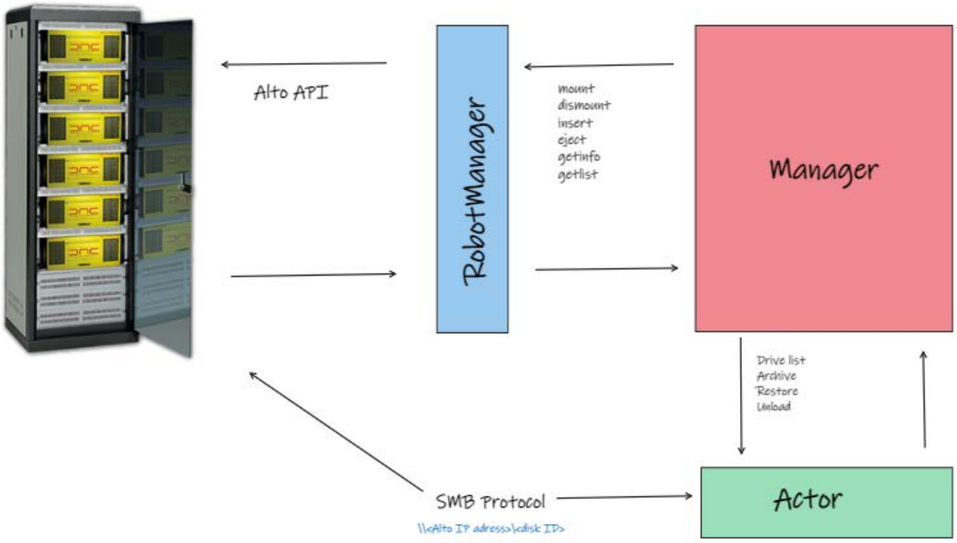

For mount operations, RobotManager needs to use the Alto API to mount the disk and make it available via SMB with these default credentials:

login: *alto*

password: *password*

The Eject command supports disk export functionality. It marks the disk as removed, so the subsequent operations on that disk are not allowed until the disk is made available again.

The Insert command is not available with Alto.

Alto Virtual Drives value can be set for an Actor from the web app Configuration > System Settings > Actors configuration page. It is on the Advanced Settings tab, in the Edit Actor page.

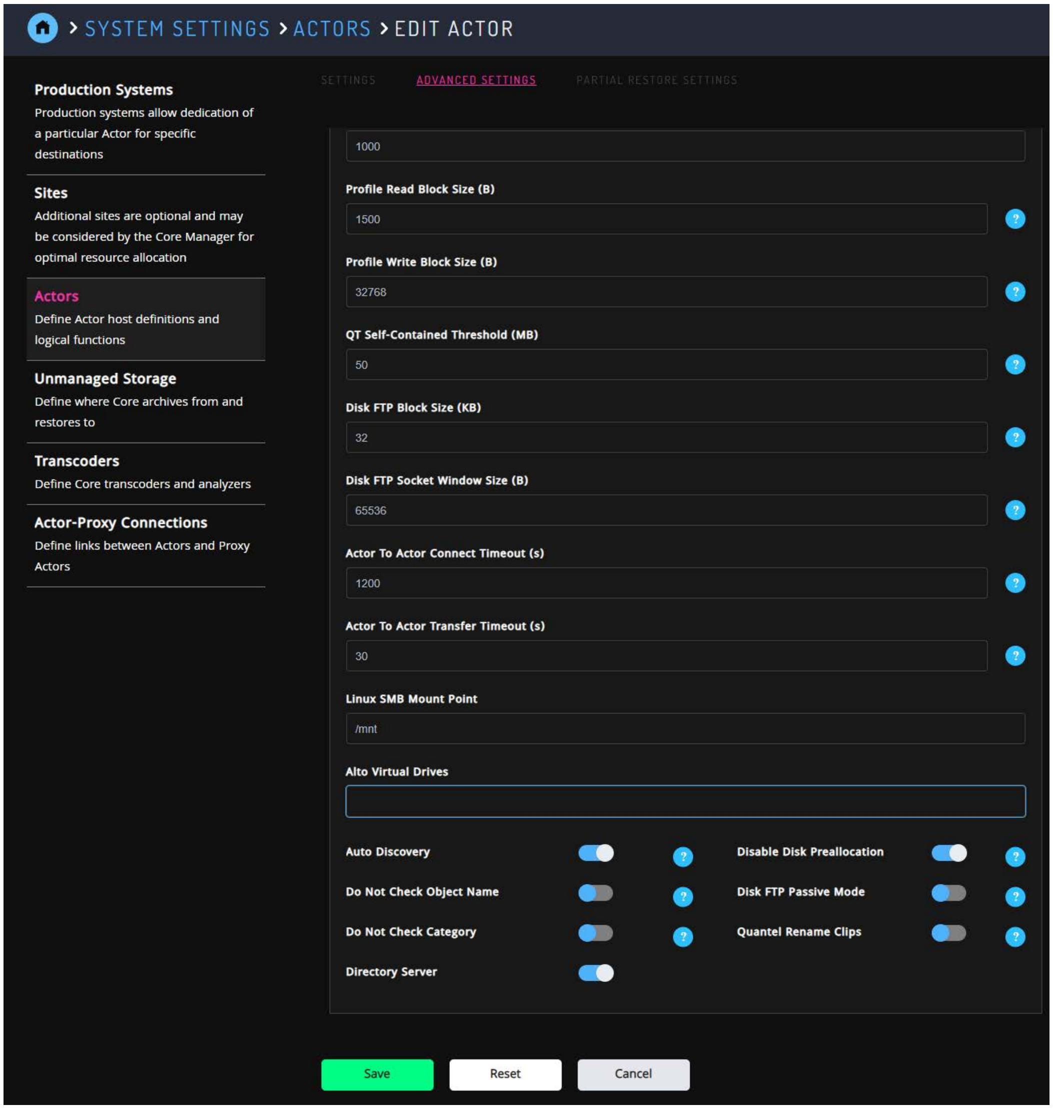

**Production Systems**
Production systems allow dedication of a particular Actor for specific destinations

**Sites**
Additional sites are optional and may be considered by the Core Manager for optimal resource allocation

**Actors**
Define Actor host definitions and logical functions

**Unmanaged Storage**
Define where Core archives from and restores to

**Transcoders**
Define Core transcoders and analyzers

**Actor-Proxy Connections**
Define links between Actors and Proxy Actors

1000

Profile Read Block Size (B): 1500

Profile Write Block Size (B): 32768

QT Self-Contained Threshold (MB): 50

Disk FTP Block Size (KB): 32

Disk FTP Socket Window Size (B): 65536

Actor To Actor Connect Timeout (s): 1200

Actor To Actor Transfer Timeout (s): 30

Linux SMB Mount Point: /mnt

Alto Virtual Drives:

| | | | |
|---|---|---|---|
| Auto Discovery | ☑ | Disable Disk Preallocation | ☑ |
| Do Not Check Object Name | ☐ | Disk FTP Passive Mode | ☐ |
| Do Not Check Category | ☐ | Quantel Rename Clips | ☐ |
| Directory Server | ☑ | | |

# Amazon S3 Integration

DIVA supports Amazon S3 AWS integration. Storage accounts allow a user to configure programmatic access to a user's AWS account. The configuration data in a DIVA storage account is exclusively used by DIVA’s Actors to query S3 storage and transfer content to and from S3 buckets.

Upon creation of a storage account, the Configuration Utility automatically creates the set of DIVA resources needed to store DIVA managed objects in S3. The resources generated by the Configuration Utility are an Array, Disk, and Actor-Disk connections.

See the DIVA Installation and Configuration Guide for more information.

## Cloudian

DIVA includes support for Cloudian integration as both Storage Accounts and Servers. Storage accounts allow users to configure programmatic access to a user's account. The configuration data in a DIVA storage account is used by the DIVA Actors to query storage and transfer content. Storage Class supported = STANDARD.

## EMC ECS (Elastic Cloud Storage) Integration

Instances stored on EMC Elastic Cloud Storage are local instances whose priority is lower than other types of local disk instances, but a higher priority than tape storage instances.

A DIVA Oracle Storage Class and Storage Location can be defined separately. If you require new cloud or local arrays in the future, you can specify all of these parameters as options. However, in DIVA both SWIFT and S3 are supported for interfacing with EMC ECS, but you cannot change the existing configuration after the Array is configured.

You can set the Media Priority of a source instance for a Restore, Partial File Restore, and Copy to Tape Group jobs, which enables you to restore an instance stored on a local non-EMC ECS array with a higher priority than an instance on an EMC ECS array. If the priorities for the media are all the same, then the DIVA decides which source instance is preferred during these jobs.

See the DIVA Installation and Configuration Guide for information.

## NetApp StorageGrid

DIVA includes support for NetApp StorageGrid integration as both Storage Accounts and Servers. Storage accounts allow users to configure programmatic access to a user's account. The configuration data in a DIVA storage account is used by the DIVA Actors to query storage and transfer content. Storage Class supported = STANDARD.

## OCI (Oracle Cloud Infrastructure)

DIVA provides support for storing your data in Oracle Cloud Infrastructure. The web app is enhanced to support OCI storage operations. OCI services combine cloud elasticity and utility with granular control, security, and predictability of your on-premise infrastructure. OCI delivers high performance, flexibility, availability, and is cost-effective.

**Note:** If you have a multiple DIVA sites, connected to the same OCI / OCI Classic storage account, you must use a different Array Name per site. The Array Name is used to uniquely identify content of an array in the cloud, and therefore must be different. This constraint is not required for other cloud vendors.

## Scality Zenko

DIVA provides support for Scality Zenko integration as both Storage Accounts and Servers. Storage accounts allow users to configure programmatic access to a user's account. The configuration data in a DIVA storage account is used by the DIVA Actors to query storage and transfer content. Storage Class supported = STANDARD.

## Data Expedition Integration

DIVA can (optionally) interface with the server named *Data Expedition Expedat Server*. The Expedat Server (also known as servedat) is very similar to the FTP_STANDARD server and CIFS, and offers AES encryption capabilities. The main difference among them is the protocol used for operations.

The Expedat Client API is integrated into the Actor computer and the Expedat server is integrated into DIVA (either on an Actor computer or other additional server within the system) just like the FTP_STANDARD server and CIFS, but is faster when used on high latency networks when using the Data Expedition Expedat MTP Protocol (a high performance file transfer protocol), which provides better bandwidth utilization.

One record is created for each Expedat Server that DIVA has to move data to or from. Although the initial feature of DIVA Connect transfer and restore is still functional in DIVA, the functionality has been enhanced and includes complex objects. With the new functionality, there are only 2 steps required for archiving objects through DIVA Connect instead of 3 steps as previously required.

### Source and Destination Server Configuration

One record is created for each Expedat server DIVA must move data from or to. Refer to the DIVA Installation and Configuration Guide for more information on Oracle Storage Cloud (OPC and OCI) and EMC integration. Here are the parameters and examples for Expedat Source and Target Servers:

- **IP Address**

  This is the IP address of the Expedat server.

  Example: 10.80.114.21

- **Source Type**

  Set this to *EXPEDAT*.

- **Connection Options**

  Hee are the connection options—ome are mandatory; others are optional:

  * `-login username`

    This is mandatory if the server is configured with authentication parameters.

    For example, *`-login moon`*.

  * `-pass password`

    This is mandatory if the server is configured with authentication parameters.

    For example, *`-pass ph4!hi4`*.

- `-port portNumber`

  This must be supplied because there is no default value.

  For example, *`-port 8080`*.
- `-license licenseCode`

  This is mandatory and is the Expedat license number.

  For example, *`-license 46FE464A98`*.
- `-encryption`

  This is optional and there are no additional parameters.

  For example, *`-encryption`*.
- `-seq_buffer_size megabytes`

  Defines the size of the Data Expedition internal buffer for each transfer. The default value is 16 MB and is sufficient for most transfers. A large buffer allows Data Expedition to continue moving data during times when the sender or receiver may not be able to process it. A small buffer consumes less memory. For example, *`-seq_buffer_size 16`*.
- `-exp_maxrate kilobytes`

  This option sets an approximate limit on the number of kilobytes per second, per transfer. The default is unlimited but can be used as an alternate method of controlling bandwidth. For example, *`-exp_maxrate 1024`*.
- `-exp_mindatagram bytes`

  This transfer protocol utilizes UDP. This option defines a minimum size for each network datagram payload that Data Expedition sends. The purpose is to prevent Data Expedition from sending too small of a packet over the network. Set this value between 2848 and 8544 when using a very fast network path (gigabit or higher) and every device along the path supports Jumbo Frames (MTU 9000). Using large datagrams can greatly reduce CPU overhead. However, using this setting without Jumbo Frames being fully supported can cause severe performance issues or loss of connectivity. For example, *`-exp_mindatagram 2848`*.

# Cloud Replicated Bucket Scanning

The purpose of this feature is to keep scanning a cloud bucket containing AXF instances (AXF and AXF_RF). This bucket is populated by a third party software:

- Can be a bucket replication software
- Can be another DIVA system

In both cases, DIVA scans the bucket for new objects and populates its database when new objects are detected.

## Supported Features

Currently supported features include the following:

- Managed disks using S3, Azure, Google Cloud and OCI are supported. Other types of disks are not supported (for example: Local, NAS, and so on).
- AXF_1.1 self-contained and AXF_RF_1.1 reference-file are supported and the format is automatically detected by the scan.
- Bad or incomplete AXF instances are currently ignored and are not rescanned.
- Complex objects are not supported yet. The service currently supports small objects (1-50 files per object).

## Starting and Stopping the Scanning Service

The scanning service is not a separate component or a Windows service; it is a functionality of DIVA. Part of this service is executed by the Manager and part is executed by an Actor selected automatically by Manager.

The service associated with the cloud disk to scan is started from the REST API (used by the web app). There is an endpoint to start the scanning service.

The service part running on the Actor only ends when the Actor is stopped if requested by the Manager. Otherwise, it keeps scanning the bucket for new AXF instances.

## Scan Events

During a scan DIVA persists events detailing the progress of the scan. The events can be retrieved using the REST API by calling *GET /disks/scans/events*. Events may be paged and filtered by the following fields:

- scanId
- serviceId
- diskName
- actorName
- objectName

- category
- description
- severity (pull-down menu)
- type (for example: AXF_SCAN, pull-down menu)
- eventType (for example: SCAN_STATISTICS, pull-down menu)
- event date-time range
- last scan date-time range

**Note:** Data that is null is not available for that event type. All times are epoch time in seconds.

## Scan Statistics

When the scan service goes through all the objects of the list, the Actor sends statistics to the Manager containing:

- The number of instances successfully detected.
- The number of bad or incomplete instances.
- The epoch time when the list started (indicates the creation/modification time of the last instance detected).
- The duration of the list scan.

To retrieve the scan status of a disk, use the `/ui/disks` endpoint. The status can be one of three values:

- NONE
- RUNNING
- STOPPED

## Rescanning Existing Objects and Instances

When the service is stopped and restarted to scan from a specific date/time, the Actor may detect instances that have been previously entered in the database. This case is covered by the Manager. The instances found by the Actor is ignored by the Manager, and there is no overwrite of the database entries.

## Rescanning Bad or Incomplete Instances

A bucket may contain bad or incomplete instances for various reasons:

- An instance is corrupted or incomplete because of a previous aborted job.
- An instance is incomplete because it has not been fully replicated yet.

The scanning service running on an Actor keeps track of the instances that failed validation during a scan.

The Actor stops testing them again after 100 retries. If there is a need to force a rescan of a bad instance beyond 100 attempts, stop the scan service first, then download, edit and overwrite a new version of the BadInstances.json file. The number of retries can be changed manually before restarting the service. The following is a sample BadInstances.json file:

```
{
  "BadInstances": [
    {
      "axf-file": "13ba7b81-9dfd-12d1-80b4-10c040243460.axf",
      "retries": 50
    },
    {
      "axf-file": "e3f4da0e-5bce-4314-817d-bf268c796c5c.axf",
      "retries": 4
    }
  ]
}
```

## Automatic Scan Restart

If the Manager or Actor service crashes during a scan, the scan restarts with the start of the corresponding service. It may take up to ten seconds following the start of the service for the scan to automatically restart.

# Arrays, Disks, and Cache

DIVA uses hard disk drive technologies for both the storage of Virtual Objects and for transient storage during data transfers (disk cache).

Any disk that DIVA uses is assigned to an array, a logical association of one or more disks for object storage. Disks that are configured as cache disks are also assigned to an array, typically named `CACHE`.

The storage of an object on a disk in DIVA is identified by the array's name rather than an individual disk itself. DIVA automatically allocates objects among two or more disks within any array.

Each disk in any array may be connected to a DIVA system either directly in an Actor host's hardware, as NAS, or connected through a SAN using Fiber Channel. In the case of SAN, it can also employ additional file system sharing software on the hosts if the disk is to be accessed simultaneously from multiple Actors.

**Note:** See the DIVA Installation and Configuration Guide for information on configuring EMC ECS components, including Object Stores in arrays, Actor-Disk connections, Array Priority, Object Storage Accounts, and validation checks.

Individual disks can be configured in an array as follows:

- **Storage Only**

  The disk is only used for storage of objects. These types of disks employ some level of RAID technology to ensure data redundancy and protection against individual hard disk failure.

- **Storage and Cache**

  The disk is used for the storage of objects and also for caching operations. Both types use separate subfolders on the disk. These types of disks employ some level of RAID technology to ensure data redundancy and protection against individual hard disk failure.

- **Cache Only**

  The disk is only used for caching, tape to tape copying, tape spanning, or tape repacking operations. These types of disks may employ RAID technology to improve performance (for example, RAID 0).

- **Storage and Nearline**

  The disk is used for the storage of objects, and also for Nearline operations. Both types use the same subfolder on the disk. These types of disks employ some level of RAID technology to ensure data redundancy and protection against individual hard disk failure.

- **Cache and Storage and Nearline**

  The disk is used for the storage of objects, caching, and Nearline operations. Both storage and Nearline types use the same subfolder on the disk. However, the cache type uses a separate subfolder. These types of disks employ some level of RAID technology to ensure data redundancy and protection against individual hard disk failure.

DIVA also enables individual disks to be configured for Read-Write access, Read-Only access, or can be disabled temporarily.

**Caution: The file system of DIVA-managed disks should never be manipulated directly by any file, DIVA, utility (such as Windows Explorer), or equivalent. If the folder structures or files are moved, renamed, or deleted, this may cause DIVA to mark that disk as Out of Order.**
**Using such a utility in any fashion completely destroys the disk's file system.**

Disks that have file sharing software installed to provide shared host access (for example, SNFS or MetaSAN) can appear as an unknown file system or as not initialized to utilities such as Windows Disk Manager.

## Disk Discovery

DIVA supports the retrieval of non-complex objects and instance metadata from an OCI/S3 cloud account to a Core database. This allows you to update an on-premise DIVA system with content from the cloud. A new web-based GUI allows a user to start, resume, or stop a disk-discovery scan and to view the status of a scan. Refer to the DIVA Installation and Configuration Guide for detailed information on how to configure Disk Discovery.

## Actor-Disk Connection Access Type

DIVA supports true READ-ONLY connections. Previously, READ-ONLY access allowed deletes. DIVA supports a separate access type called READ/DELETE to allow both operations. Additionally, READ/WRITE has been re-labeled as READ/WRITE/DELETE because it supports all three options.

**Note:** If a disk is currently under scan, changing the associated connection's access type to Read/Write/Delete still won't allow writes to the disk. Only after the scan is stopped will DIVA allow writes if the access type for that disk permits it.

The Cloud Disks page in Resource Management for the DIVA web app displays the cloud disks and their scanning status in the main table:

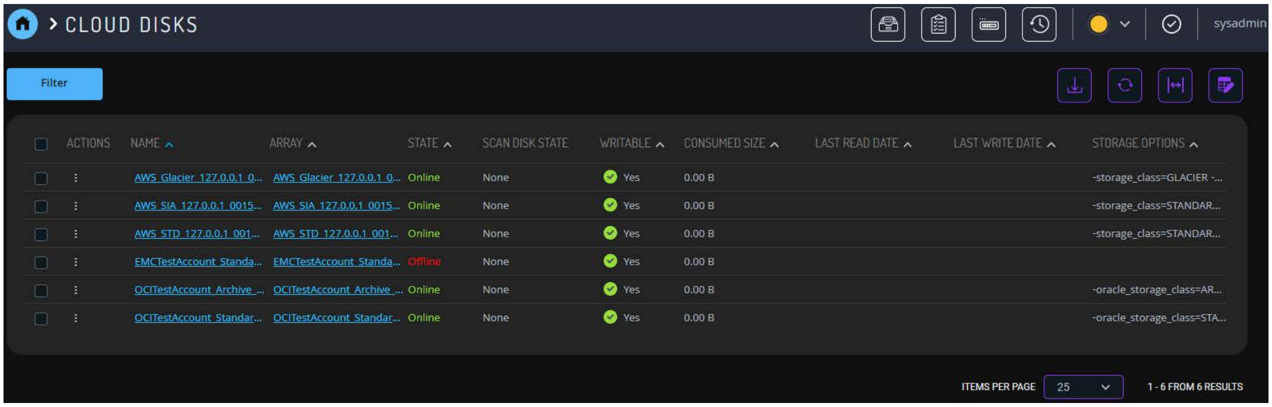

| □ | ACTIONS | NAME | ARRAY | STATE | SCAN DISK STATE | WRITABLE | CONSUMED SIZE | LAST READ DATE | LAST WRITE DATE | STORAGE OPTIONS |
|---|---|---|---|---|---|---|---|---|---|---|
| □ | ⋮ | AWS_Glacier_127.0.0.1_0... | AWS_Glacier_127.0.0.1_0... | Online | None | ✓ Yes | 0.00 B | | | -storage_class=GLACIER -... |
| □ | ⋮ | AWS_SIA_127.0.0.1_0015... | AWS_SIA_127.0.0.1_0015... | Online | None | ✓ Yes | 0.00 B | | | -storage_class=STANDAR... |
| □ | ⋮ | AWS_STD_127.0.0.1_001... | AWS_STD_127.0.0.1_001... | Online | None | ✓ Yes | 0.00 B | | | -storage_class=STANDAR... |
| □ | ⋮ | EMCTestAccount_Standa... | EMCTestAccount_Standa... | Offline | None | ✓ Yes | 0.00 B | | | |
| □ | ⋮ | OCITestAccount_Archive_... | OCITestAccount_Archive_... | Online | None | ✓ Yes | 0.00 B | | | -oracle_storage_class=AR... |
| □ | ⋮ | OCITestAccount_Standar... | OCITestAccount_Standar... | Online | None | ✓ Yes | 0.00 B | | | -oracle_storage_class=STA... |

ITEMS PER PAGE 25 1 - 6 FROM 6 RESULTS

Start Cloud Bucket Scanning from the Action menu for an individual disk. When a Cloud Disk is running a scan, the Action menu also allows you to stop the scan.

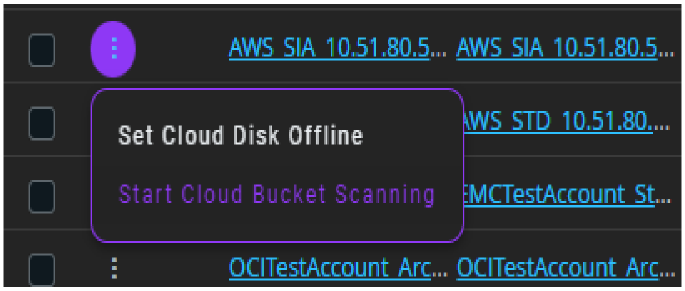
AWS_SIA_10.51.80.5... AWS_SIA_10.51.80.5...
Set Cloud Disk Offline
Start Cloud Bucket Scanning
AWS_STD_10.51.80....
EMCTestAccount_St...
OCITestAccount_Arc... OCITestAccount_Arc...

DIVA displays this form to enter the scan settings:

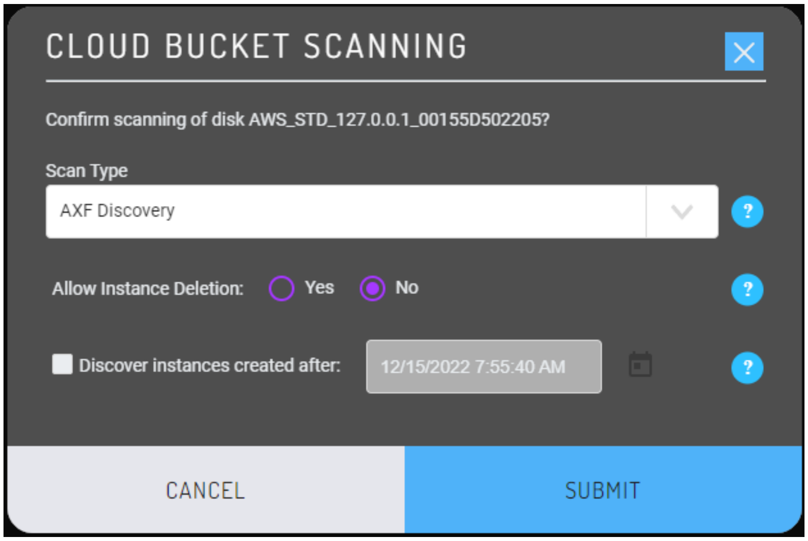

CLOUD BUCKET SCANNING

Confirm scanning of disk AWS_STD_127.0.0.1_00155D502205?

Scan Type: AXF Discovery

Allow Instance Deletion: ○ Yes ● No

□ Discover instances created after: 12/15/2022 7:55:40 AM

The Cloud Disk Properties page displays the scan events associated with the disk in a table at the bottom of the page. Events can be filtered by Severity, Type, Actor, Object, or Description. You can also control the scan state from the property page in the Actions pull-down menu:

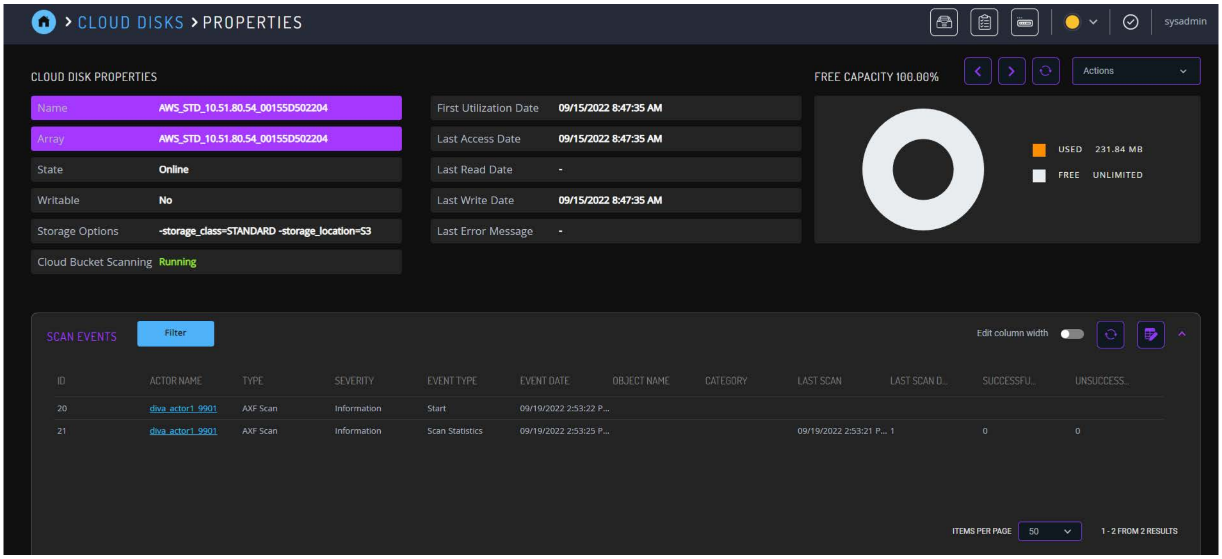

CLOUD DISK PROPERTIES

| | |
|---|---|
| Name | AWS_STD_10.51.80.54_00155D502204 |
| Array | AWS_STD_10.51.80.54_00155D502204 |
| State | Online |
| Writable | No |
| Storage Options | -storage_class=STANDARD -storage_location=S3 |
| Cloud Bucket Scanning | Running |

| | |
|---|---|
| First Utilization Date | 09/15/2022 8:47:35 AM |
| Last Access Date | 09/15/2022 8:47:35 AM |
| Last Read Date | - |
| Last Write Date | 09/15/2022 8:47:35 AM |
| Last Error Message | - |

FREE CAPACITY 100.00%

| | | |
|---|---|---|
| 🟧 | USED | 231.84 MB |
| ⬜ | FREE | UNLIMITED |

SCAN EVENTS

| ID | ACTOR NAME | TYPE | SEVERITY | EVENT TYPE | EVENT DATE | OBJECT NAME | CATEGORY | LAST SCAN | LAST SCAN D... | SUCCESSFU... | UNSUCCESS... |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 20 | diva_actor1_9901 | AXF Scan | Information | Start | 09/19/2022 2:53:22 P... | | | | | | |
| 21 | diva_actor1_9901 | AXF Scan | Information | Scan Statistics | 09/19/2022 2:53:25 P... | | | 09/19/2022 2:53:21 P... | 1 | 0 | 0 |

ITEMS PER PAGE 50 1 - 2 FROM 2 RESULTS

# Tape Groups and Sets

Disks are logically assigned to arrays for the storage of objects, but tapes are logically associated together in Tape Groups.

Tapes are initially divided into Sets, and assigned a number called a Set ID. A Set ID enables partitioning of pools of tapes in a library and assigning them for use with specific Tape Groups. A Tape Group draws upon the pools by associating the Tape Group with a Set ID.

More than one Tape Group may use the same Set ID. An unused tape does not actually belong to any of those Tape Groups until DIVA writes the first object to that tape. When all objects have been deleted from a Tape Group-assigned tape, it is unassigned from that Tape Group, and may subsequently be assigned to another Tape Group using the same Set ID.

Since Tape Groups are a user-defined concept they can differ from one DIVA system to another. The exception is the Default Tape Group, which exists in all installations and can neither be renamed nor removed. In a DIVA system, Tape Groups are created and managed on the web app on the Configuration > Library Storage > Tape Groups page.

When a tape is assigned a Set ID of 99 it indicates to DIVA that the tape is not to be used and is not related to the operation of DIVA. Examples are tapes that belong to a non-DIVA application in a shared library environment or the library's cleaning tapes.

This figure illustrates how tape sets and Tape Groups are associated and used:

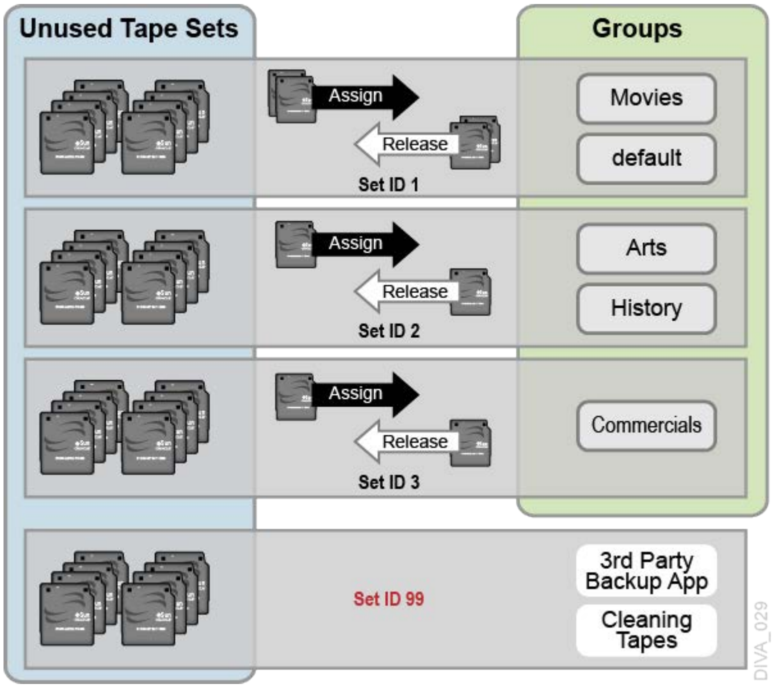

## Tape Compression

Tape compression is supported at the Tape Group level. Tape compression is enabled or disabled in the web app on the Configuration > Library Storage > Tape Groups > Edit Tape Groups page.

When tape compression is enabled, any empty tape assigned to the Tape Group has compression enabled, and instances written to the tape are compressed. Tapes assigned to the Tape Group before compression was enabled remain uncompressed, and instances written to the uncompressed tape are uncompressed.

When exporting a tape, compression is tracked using the `isCompressionEnabled` attribute. This attribute value can be either true or false.

When writing to a Tape Group with a compression setting (enabled or disabled), only used tapes with the same compression setting, or empty tapes associated with the set linked to the Tape Group, are used.

## Tape Group Encryption

Tape drive encryption securely supports bulk tape migration between DIVA systems.

See the DIVA Installation and Configuration Guide for detailed configuration information.

## Sony ODA Drives

DIVA supports the Sony optical drives and their WORM media (using a UDS format). Only AXF formatted objects can be written to Blu-ray discs. The drives are controlled by the Robot DIVA and the media is viewed as a tape cartridge.

These drives are shown as Unknown Medium Changer under the Medium Changer devices section in the Windows Device Manager because there are no device drivers for them. The drive itself also displays as an Optical SCSI Device with the make and model number under the Disk Drives section.

### Using Optical Drives and Discs

The following list is additional information related to the use of the optical drives and discs:

- In the web app the optical discs are displayed under the Drives tab.
- Write-Once media must be finalized, and therefore zero remaining space is reported to DIVA.
- Objects are spanned when there is 100 MB remaining on the disc so that there is enough space left to finalize the disc. After an object is spanned, the disc is considered full and is automatically finalized.

- The Actor auto-finalizes the disks when there is 500 MB of space remaining (unless an object was spanned). However a disc can be finalized manually through the Optical Disc Archive Utility.
- If a drive is mounted manually and viewed in Windows Explorer, the numeric value at the beginning of each object's file name identifies the object's location on tape.

## Tape Spanning

When the capacity of a Tape Group begins to reach full capacity (that is, the Tape Group's associated Set ID has no more blank tapes to draw upon), DIVA may attempt to maximize the storage utilization of the existing tapes in the Tape Group by filling the remaining free space of each tape by segmenting the object across two or more tapes (tape spanning).

By default, the tape spanning feature is configured in the DIVA Configuration file to *not* span tapes. If an object cannot be spanned across the remaining free space of two tapes within that Tape Group, the job is terminated by DIVA. Tape spanning can be configured to span across more than two tapes at the site, or disabled altogether in the diva.conf configuration file.

During the restore of a spanned object, DIVA mounts all associated spanned tapes and automatically joins the spanned object back together. It cannot do this directly, and must copy all segments of the spanned file to a cache disk first. Therefore, restoring a spanned object must use a Cache Only or Cache and Direct QOS. A Direct QOS results in the job terminating.

For Write-Once media, objects are spanned when there is 100 MB remaining so that there is space left and the disc can be finalized. After an object is spanned the disc is considered full and is automatically finalized.

If spanning is disabled and an object is to large to fit on the selected tape, DIVA retries to fit all content on a single empty tape. The retry logic can be configured to optionally retry with a used tape with the least or most remaining space. The type of tape DIVA selects on retry is configurable in the DIVA configuration file through the `DIVA_RETRY_ON_SPAN_REJECTED_ALGORITHM` configuration parameter.

**Note:** Any job that writes to a tape linked to a clone, terminates on a span. In addition, any attempt to clone a tape with spanned instances aborts.

## Protected Mode

When a tape is ejected from the library it is automatically set to Protected Mode. When this attribute is set it prevents further archive operations from being performed on the tape and prevents the tape from being repacked.

DIVA assumes that when a previously ejected tape is reinserted into a library to perform a restore operation it subsequently is ejected and put back into offline storage. Without the Protected Mode feature new Virtual Objects may get written to the tape while it is temporarily in the library and prevent it from being ejected without first moving these required objects to another tape.

Write operations on a protected tape are not allowed unless the protected attribute is set back to false in the web app after the tape is reinserted into the library. This attribute does not prevent delete operations on instances located on these tapes (whether internalized or externalized).

This attribute may also need to be reset on a tape if the tape was mistakenly ejected from a library, or if the tape was stuck in a tape drive and removed by opening the library door and manually ejected from the library. When the library is then resynchronized with the Core database, the missing tape is deemed externalized and Protected Mode set to true (the tape is protected).

## Tape Label Management

When a tape is first mounted and objects written to it, DIVA writes a label to the beginning of that tape. The label contains important information relating to the management of objects written to, or deleted from, the tape during archive operations. From an operational perspective the most important information in the tape label is the barcode number of that tape. The barcode is an alphanumeric number on the physical label adhered to the back of the cartridge, and is also written to the label on the magnetic media of the tape.

DIVA automatically checks the label written on tape whenever a tape is mounted to verify that it matches the tape barcode label it instructed the tape library to mount.

This mechanism provides the following two safety features:

- Confirmation that the mapping between the physical drives in the library matches that of the logical connections to each tape drive from the Actor. This prevents data being written to the wrong tape if there is a configuration mismatch between the physical drives in the library.
- Prevents tapes with foreign labels (that is, tapes previously used by another archive system) from being overwritten in error. This behavior is for environments where DIVA shares a library with another archive application and the tapes used by that archive application have not been set to Set ID 99.

If DIVA identifies a mismatch between the expected label and that of the tape, it generates an I/O Label Error, and the tape is set to Not Writable and isn’t selected for any further write operations.

# Media Storage Formats

This section discusses the media formats available in DIVA.

## AXF Disk and Tape Storage Formats

AXF (Archive eXchange Format) is an open format that supports interoperability among disparate content storage systems and ensures the content's long-term availability regardless of how storage or file system technology evolves.

An AXF Virtual Object is an IT-centric file container that can encapsulate any number, and any type, of files in a fully self-contained and self-describing package. The encapsulated package contains its own internal file system, which shields your valuable data from the underlying operating system and storage technology. It's like a file system within a file that can store any type of data on any type of storage media.

### Native File and Folder Support

Users can see their files and folders in native format on archive devices rather than as an AXF container file. You can also access files and folders on storage devices like object storage. This access opens the archive to the use of third party software to perform operations on the archive (for example, metadata collection, face recognition, transcoding, and so on).

The following list identifies the different AXF formats available:

- **LTFS_AXF_1.1**

  This new format offers the same features as AXF_1.1 on an LTFS formatted tape and is only supported on tape. This format is not recommended for complex objects because it would generate very large LTFS indexes.

  The tape block size must be set to 524288 Bytes or greater; LTFS does not support lower block sizes.

  If an LTFS tape is loaded by LTFS software on a standalone drive, the contents of that tape can be accessed using Windows Explorer.

**WARNING: Accessing LTFS tape content in Windows Explorer should only be used for recovery purposes, and the LTFS software must not be running concurrently with DIVA on the same drives.**

The following features are supported by LTFS AXF:

- Spanning
- Drive compression
- Drive encryption is supported, but not recommended if there is no procedure to export and load encryption keys to the LTFS software.

- **AXF_RF_1.1**

  This format uses the AXF 1.1 structures, but AXF files won't contain any overhead. This format is supported on disk arrays and object storage only.

**Note:** Telestream does not recommend using AXF_RF_1.1 format with complex objects.

  This format allows a user to see the files of an object on disk or cloud. When archiving complex objects with small files, performances are better using AXF_1.1 because the files are wrapped into an AXF container (this is the reason for the previous the note). When this format is chosen for a cloud storage array, there are limitations on the server side to the size of file that can be created. This limitation is approximately 5TB per file. Telestream recommends that you use AXF_1.1 instead of AXF_RF_1.1 if a DIVA Object contains a file larger than 5TB.

- **AXF_1.1**

  This format is compliant with AXF 1.1 standards and is the recommended format when archiving complex objects. DIVA creates a new AXF segment every 500GB to avoid compromising multi-part upload performance when this format is selected for cloud storage. Some object storage vendors limit the number of parts per object to 10000. The 500GB AXF segment limit maintains a reasonable size for each part.

- **AXF_1.0**

  This format is compliant with AXF 1.0 standards.

- **AXF**

  This is redirected to AXF_1.1.

- **LEGACY**

  This is the formal archive format used by DIVA (index.txt, 00000001, 00000002, and so on)

## Storage Media Format

In DIVA, a Tape Group or disk array has a Media Format parameter that indicates which storage media format to use when creating new archived objects. The Media Format can be set to either DIVA Legacy format or one of the AXF formats. This setting can be changed at any time and it does not influence content already stored. Therefore, it is possible to have more than one storage media format within Tape Groups and Disk Arrays.

DIVA writes an object instance in one, and only one, media format. Therefore, if an object spans tapes, each tape used as part of an object instance is written in the same media format. In DIVA, an object can contain multiple instances, each of which can be stored in either Legacy or one of the AXF formats.

Complex objects must be stored in AXF format. Because all complex objects are written in the AXF format, every instance of a complex object is in the AXF format.

## Tape Storage Media Format

Although a Tape Group can contain more than one storage format, an individual tape has (at most) one storage media format. DIVA assigns the tape media format to an empty tape when it writes the first object to that tape. The tape is assigned the format of the Tape Group that appears in the job. After the media format for a tape is assigned, it cannot be changed unless all objects on the tape are deleted. After deletion of all objects from a tape, the tape's format becomes unassigned until DIVA writes content to the tape again. If the tape was in use, the tape format cannot change unless it is empty and cleared.

Both Legacy and AXF formatted tapes can exist in the same Tape Group. Objects in AXF format are only written to AXF formatted tapes, and objects in Legacy format are only written to Legacy formatted tapes even though they are in the same Tape Group.

A Repack job always writes the destination tape in the same media format as the source tape. Similarly, tape spanning operations always uses the same format across all tapes storing spanned objects.

## Disk Storage Media Format

Unlike tapes, disks do not have a format. DIVA allows storing objects in different media formats on the same disk. If a disk contains objects in Legacy format and that disk is then assigned to an AXF formatted array, it still contains objects in Legacy format. However, new objects written to the disk are written in AXF format.

## Object Instances Media Format

Every tape and disk object instance is assigned a format of Legacy or AXF. The format of a tape or disk instance is assigned when the instance is created and is the format of the tape on which the instance resides. All instances on a tape must have the same format.

If a disk instance is non-complex and permanent (not a cache instance) it is stored in the format of the destination array. If a cache instance is non-complex it is stored in the format of the Tape Group specified in the job.

Tape Groups or arrays used by complex object jobs must be in an AXF format because complex objects cannot be stored in Legacy format. Therefore, any instance of a complex object is in the AXF format.

A migration job must be used to change a tape format from Legacy to AXF: repacking a tape does not change the tape format. Repacking of existing Legacy format objects retains the format of the tape even if the Tape Group format was updated in the configuration from Legacy to AXF.

# Objects

Each asset that is archived to DIVA is called an object.

An object is a DIVA logical container for all files consisting of an asset from the original source. Assets from some sources may have separate video, audio and metadata files. When archived in DIVA all of these files are referenced as a single object. When the object is restored to a destination, all files that were originally associated with that asset are automatically restored by DIVA.

An object is uniquely identified in DIVA by its name and Collection. The Object Name does not necessarily need to match that of the source file being archived. DIVA always restores the archived files as they were archived, regardless of the DIVA Object Name. Therefore, the same source file can be archived more than once in the same Collection, if each instance has a unique Object Name.

After an object exists within DIVA, it cannot be replaced unless it is first deleted. If an Archive job uses the same name and Collection of an existing object, DIVA automatically stops the job. However, multiple copies (or instances) of an object can be created after the asset is archived.

If a source asset is to be stored in a variety of encoding formats (for example, MPEG2 Long GOP, DV50, or low resolution proxies), you can use specific categories to archive the same object based on its encoding format.

## Complex Objects

When the Metadata Database feature is enabled, the complex object feature is available. DIVA can track significantly more than the 10,000 file per object limit set for non-complex objects using complex objects. The actual amount scales with system processing power and storage capacity. A complex object stores more information about the files and folders in an archive, such as subtotals for each directory.

When an object is archived, DIVA determines whether the new object should be complex or non-complex based on its number of components (files). If the number of components is greater than 5,000 (the default complex object threshold, which is configurable), the object becomes a complex object: otherwise, the object is non-complex. When an object is deemed a complex object, it is always be complex— even if it is copied using the Copy As command or imported using the Export/Import Utility.

## Complex versus Non-Complex Objects

A complex object differs from a non-complex object in some key ways. For example, the file and folder metadata information of a complex object is stored in the *Metadata Database (MDS)* not in the Postgres database. The file contains the file names, folder names, checksums, and files sizes. The directory that contains these files is the Metadata Database Root Directory (the following section describes how to configure this parameter). A complex object must be stored in AXF format either on tape or on disk.

A complex object can contain hundreds of thousands of files. In the web app the entire set of files on a tape are not displayed in the Object Properties and Tapes dialog boxes—only a single placeholder file is shown to represent the complex object.

Certain API operations used in Avid Connectivity (such as *GetByFilename* and *DeleteByFilename*) are not supported for complex objects.

A complex object maintains information about the folders and files in the archive. Complex objects store subtotals for each folder, including the total number of files and subdirectories within the folder, and the total size of all files within the folder and any subfolders.

The Complex Object Threshold is a configurable parameter used by DIVA to determine whether a new object should be complex. If a new object has many components (files) that exceeds the threshold, the object automatically becomes a complex object. This value is set in the diva.conf configuration file. Telestream recommends leaving the threshold at the default value (5000 components) unless there is a specific reason to adjust the value.

Not all operations are supported for complex objects: The Delete on Source feature is disabled and the checksum features Verify on Archive and Verify on Restore are also disabled. DIVA Connect does not support replication of complex objects.

## Complex Objects and FTP

When archiving Complex objects with the FTP protocol and using FileZilla with default settings, the transfer typically fails when archiving any object with more than approximately 3900 files. There are two reasons for this possible failure:

- The Actor connection times out before the size of the object can be computed.
- A job stops in the middle of the transfer because the FTP server (for example a FileZilla server) is consuming all of the available sockets.

**Note:** Telestream only supports Linux-based FTP servers on DIVA systems running in the Linux environment, not FileZilla and IIS FTP servers.

Actor connection timeouts can be resolved by setting the following two parameters either in the Server Command Options, or in the options of the command itself as follows:

```
-transfer_timeout 1200
-list_timeout 600
```

Telestream also recommends setting the corresponding parameters in the FileZilla server under the General Settings:

```
Connections Timeout = 600
No Transfer Timeout = 1200 (this is the default)
```

If a termination occurs, which may happen during transfer, there are two registry parameters that must be created or modified (typically created):

```
TcpTimedWaitDelay = 10
MaxUserPort = 90000
```

Contact Telestream (see *Telestream Contact Information*) for more information on these parameters and to make FTP server and computer registry changes if no qualified personnel are on site.

## DIVA Connect Complex Object WAN Transfers

DIVA provides optional WAN acceleration that enables it to take full advantage of long distance, high latency, network paths (such as private site to site links or the public Internet), and can perform transfers of complex objects efficiently using Data Expedition MTP/IP protocol.

### Example

1. DIVA instance 1 restores the complex object to the DIVA 2 system by first creating a new AXF file on the DIVA 2 system's Data Expedition server.
2. DIVA 1 restores all of the files from the local storage to the new AXF file created on the DIVA 2 Data Expedition server.
3. The DIVA 2 system creates a new AXF file on the destination (tape, disk, and so on).
4. The DIVA 2 system archives all of the files from the Data Expedition AXF file (created by DIVA 1 on the Data Expedition server) to the newly created AXF file on the destination.

See the DIVA Connect Installation, Configuration, and Operations Guide (https://www.telestream.net/telestream-support/) or contact Telestream (see *Telestream Contact Information*) for more information and assistance.

## Object Instances

The storage managed by DIVA falls into three distinct categories:

- Online Storage (tapes within a library)
- Nearline Storage (disks and cloud storage)
- Offline Storage (externalized tapes)

The name and Collection of an object in DIVA must be unique. However, multiple copies of that object can be created in one or all three of the above classes. Each copy of an object (including the original archived object itself) is known as an Object Instance.

Apart from creating backup copies, the object instances concept also allows life cycling of material within DIVA. A object may initially be created in online storage for rapid access and also backup instances created on one or more tapes. When the object is no longer required for online or Nearline access, the disk instance can be deleted and the tape externalized. Automatic life cycling of objects, based on their age and location within the archive, can be provided by the SPM (Storage Policy Manager) option.

The first instance of an object is created when it is first archived to DIVA. Additional instances of the archived object can then be created with the Copy command.

An additional instance of an object cannot be created by re-archiving the original object with the same name and Collection. This job is automatically stopped by DIVA with the "Object already exists within DIVA Core" error.

Instances are initially numbered sequentially with the original object that is archived to DIVA being Instance 0. As new instances are created and older instances deleted, the instance numbering may no longer be sequential when an object's properties are viewed in the web app on the Content Management > Catalog Browsing Page. However, an instance number from a previously deleted instance may be subsequently reused by DIVA in additional copy jobs.

The following restrictions apply to creating new instances of an object in DIVA:

- A Tape Group can contain two instances of the same object if both can be located on separate tapes. If no additional tapes for that Tape Group are available to store the second instance the copy job is terminated.
- A disk array can contain two instances of the same object if both can be located on separate disks within the array. If no additional disk is available, the copy job is terminated.

When an object has multiple instances within the archive and a restore job is issued, DIVA performs as follows:

- If no instance number is specified in the job, DIVA chooses the instance that allows the job to complete in the shortest possible time. A disk instance is preferred over a tape instance. However, a tape instance may be selected in some configurations if the QOS specified in the job is Cache Only or Cache and Direct.
- If no instance number is specified in the restore job and a disk instance exists, but the disk is offline, the tape instance is automatically selected.
- If two or more instances are present on tape and no disk instances exist, and one tape is currently in use by another job (or is externalized), the tape containing the other instance is automatically selected.
- If two or more instances exist on tape, and a read error occurs on the first instance selected, the job is automatically attempted on the other instances until the job is completed successfully. If no instances can be read, the job is terminated.
- If a specific Instance Number is specified in the restore job, DIVA only uses that instance. If the media containing the instance is offline (for disks), externalized (for tapes), or an I/O or read error occurs, the job is terminated.

## Metadata Database (MDS)

To effectively operate with large volumes of files and folders and other metadata, DIVA stores the metadata separately from the Postgres database in the Core Metadata Database (*MongoDB*). The Core metadata database contains files stored in a file system local to the DIVA. The directory that contains these files is the Metadata Database Root Folder.

The metadata database has very high performance and almost unlimited scalability. The metadata database should be treated with the same caution as the Postgres database, and should be backed up at regular intervals using the DIVA Backup Service. The metadata Database is backed by MongoDB.

# Jobs

A job is a command that is issued to DIVA to perform an operation. Jobs can be issued through the web app or an Archive Initiator.

The most common job types are for transferring content to the archive (referred to as an Archive job), or transferring content out of the archive (referred to as a Restore or Partial File Restore job).

You use other job types for managing the objects within the archive after they have been created. Examples of other job types are Copy, Delete, and Repack Tape jobs.

Each job is automatically given a unique identifier by DIVA (called the Job ID), which can be used later to retrieve event logs or other properties of each job. DIVA stores the records of up to 50,000 jobs in its database.

Because multiple jobs may be received simultaneously by DIVA, they are all placed into a queue and are executed on a first come, first served basis. The execution order of jobs can be prioritized using the Job Priority parameter. The Content Management > Jobs page in the web app displays the queue of jobs that are currently being processed by DIVA. The Content Management > Job History page displays previously executed jobs and the status of those jobs.

When restoring the same file to the same destination twice in parallel, the behavior on Windows and Linux is different. On Windows, the first restore (they cannot arrive exactly at the same time) locks the file so that the second one terminates. On Linux, there is no such lock at the file system level. Both restores are executed at the same time, and both write to the same file. The content of the resultant file is not predictable.

The DIVA available Job Options are as follows:

- Archive Jobs

  ```
  -delete_on_source
  ```

- Restore Jobs

  ```
  -do_not_overwrite
  -do_not_check_existence
  -delete_and_write
  ```

Job Options take precedence over the normal Additional Service specification. Also, the normal Additional Service specification takes precedence over the Server Connect Options.

The Additional Services available for a Restore job can also be specified in the Server Connect options. If specified, the server uses the Additional Service setting as a default. You can over ride this by specifying the Additional Service at a job level in the normal way, or as a new Job Option. Because these connect options are specific to a Restore job, the options are ignored for any other type of jobs using the server.

# Job Types

This topic describes the DIVA job types.

When connected to DIVA, access the Content Management > Catalog Browsing page in the web app to execute jobs to be issued to DIVA. Click the three dots next to the desired object to display the context menu with the various job types for that object. A third party initiator application (for example, an Automation System) can be used instead of, or in addition to, the web app interface. The available options in the context menu redetermined by the logged in user's assigned role in the user profile.

The different options available from an object's context menu are as follows:

- **Search in Job History**

  Displays the job history for the selected object.

- **Assign Storage Plan**

  Enable assigning the object to a SPM storage plan.

- **Restore Object**

  Displays a form to configure copying a file from DIVA to a single target.

- **Partial Restore**

  Displays a form to select which file, or files, specifically to copy from DIVA to a Target Server.

- **Multiple Restore**

  Displays a form to configure restoring an object from DIVA to more than one target simultaneously.

- **Delete Object**

  Displays a form to configure deleting all instances, or a selected instance, of an object.

- **Copy**

  Displays a form to configure copying an existing object to another Tape Group.

- **Copy As**

  Displays a form to configure copying an existing object to another name, Tape Group, or Collection.

- **Stage**

  Displays a form to configure staging content in the cloud so that it is readily available when Copy and Restore jobs are submitted for the same content.

# Amazon S3 Transfers

## Archiving to an S3 Disk

DIVA enables you to archive content from and to a regular source such as an FTP server to and from an Amazon S3 disk. The following figure displays the Archive Object screen for transfers. Fill in all required settings, select the Source and Media (in this case an S3 disk), and click Submit to execute the job using the configuration identified:

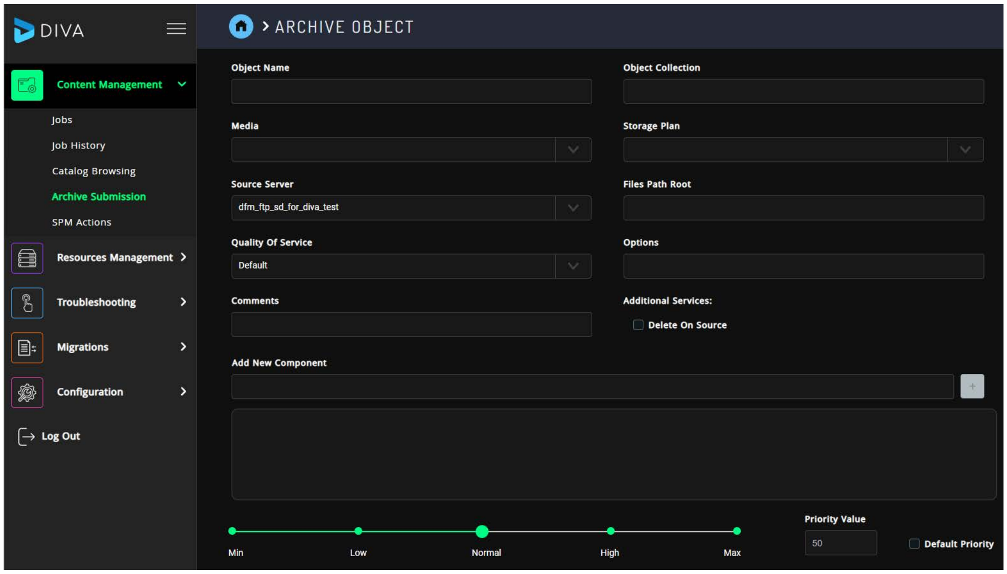

Object Name

Object Collection

Media

Storage Plan

Source Server: dfm_ftp_sd_for_diva_test

Files Path Root

Quality Of Service: Default

Options

Comments

Additional Services:

□ Delete On Source

Add New Component

Priority Value: 50

□ Default Priority

Files in a folder of an FTP server are split into 5 AXF segments (determined by number of Threads Per Transfer specified in the Storage Accounts configuration) and transferred to an Amazon S3 bucket using 5 threads within the Actor. DIVA automatically creates the bucket named `diva-<unique bucket id>-<region>-<index>` where the index increases every 100,000 instances. This is specified in the Max Instances Per Bucket setting of the array's configuration. The unique bucket id is generated on creation of the storage account.

### Restoring from an S3 Disk

Users can restore an object on an Amazon S3 disk to a non-S3 destination, for example an FTP server. Users can also restore an object on an Amazon S3 disk as an AXF file to an S3 destination. In the web app navigate to the Content Management > Catalog Browsing page and click the three dots next to the object to be restored. From the context menu select the type of restore desired and fill in the form information, then click Submit to execute the job.

### Copying from or to an S3 Disk

Users can copy an object from one S3 disk to another S3 disk in the same AWS account, or in another AWS account. It is also possible to copy from an OCI Disk in an Oracle Object Storage Account to an Amazon S3 Disk in an Amazon account.

## Oracle Storage Cloud Transfers

The Oracle Storage Cloud is an object storage solution that offers two types of accounts usable with DIVA: metered and non-metered accounts. Visit http://docs.oracle.com/cd/

E60880_01/VLPFN/whatis.htm#BABDADAE for information on Oracle Storage Cloud storage accounts.

The non-metered account enables the creation of standard class containers. Virtual Objects written inside standard containers are accessible immediately at any time.

With a metered account, DIVA can archive to standard and also to archive class containers. With archive containers, objects written to a deep archive storage device require a restoration process before they can be downloaded.

An object located in deep archive requires a maximum of 4 hours to restore to a configured Target Server because the content is first transferred from tape to Cloud cache, and then transferred from cache to the final destination.

When a Restore job is created for an object with a cloud instance, DIVA always attempts to restore a local (non-cloud) instance of a Virtual Object. If all local instances are offline, no local instances exist, or when a cloud instance is explicitly requested (a Restore Instance job), then DIVA restores from a cloud instance.

Only Actors configured for CLOUD ARCHIVE can transfer content to the cloud. Only Actors configured for CLOUD RESTORE can transfer content from the cloud.

## Object Storage Destinations

DIVA enables restoring content to a destination, and archiving content from a source, linked to an Oracle Object Storage account. Any type of object can be restored to these destinations. However, these destinations do not support symbolic links.

The Files Path Root for the destination must contain a value, and can contain an optional prefix. The value identifies the name for the target container. Use the optional prefix if the object is not to be restored to the container root directory. The prefix must be separated from the container name using either / or \. For example, `container`, `container\folder`, and `container/subdir1/subdir2` are all valid paths.

## EMC ECS Object Store Integration

DIVA supports local arrays that include disks with Swift interfaces—for example, an EMC ECS Object Store.

During an upgrade from an earlier DIVA version, all disk instances with an `ARCHIVE` or `STANDARD` Storage Class are updated with a storage option containing both of the following:

- `-storage_location=CLOUD`
- `-oracle_storage_class={ARCHIVE|STANDARD}`

All disk instances with a `NONE` Storage Class are updated with a storage option containing both of the following:

- `-storage_location=LOCAL`
- `-oracle_storage_class=NONE`

All Actor-Disk connections with cloud as the interface are updated to Swift for the interface.

See the DIVA Installation and Configuration Guide for detailed configuration information.

## True Remaining Tape Size and Last Written Position

For some specific tape drives (Oracle T10K and IBM LTO) the Actor now returns the True Remaining Size on the tape and the Last Written Position on tape to DIVA during a transfer of content to tape. The remaining size is given in number of bytes of uncompressed data.

DIVA uses the remaining size and last written position (instead of relying on the size of the tape type) to obtain the true total and remaining size on the tape in all tape based operations.

Export and import operations also now include the total tape size.

## Archive Submissions

An Archive operation is defined as the transfer of files to DIVA. The files are then stored as an object. Archive jobs are issued by navigating to Content Management > Archive Submission in the web app. The form submits an Object Archive job to the Manager for processing.

The following fields are included in the Archive Submission page:

- **Object Name**

  The name of the object to be archived.

- **Object Collection**

  The Collection of the object to be archived.

- **Media**

  This field designates either a group of tapes or an array of disks declared in the configuration where the instance must be created. When this parameter is a null string, the default Tape Group of tapes named DEFAULT is used.

- **Storage Plan**

  This field defines the Storage Plan to use for this Virtual Object. If no Storage Plan is assigned the default Storage Plan is used.

- **Source Server**

  The name of the source (for example, a video server, browsing server, and so on). This name must be known to the DIVA configuration.

- **Files Path Root**

  The root folder for the files (see the examples in the following section).

- **Quality of Service**

  One of the following codes (see *Quality of Service* for detailed descriptions):

  - `DIVA_QOS_DEFAULT`

    Archiving is performed according to the default Quality of Service (currently direct and cache for archive operations).

  - `DIVA_QOS_CACHE_ONLY`

    Use cache archive only.

  - `DIVA_QOS_DIRECT_ONLY`

    Use direct archive only—no disk instance is created.

  - `DIVA_QOS_CACHE_AND_DIRECT`

    Use cache archive if available, or direct archive if cache archive is not available.

  - `DIVA_QOS_DIRECT_AND_CACHE`

    Use direct archive if available, or cache archive if direct archive is not available.

    Additional and optional services are available. To request those services, use a logical `OR` between the previously documented Quality of Service parameter and the following constant:

  - `DIVA_ARCHIVE_SERVICE_DELETE_ON_SOURCE`

    Delete source files when the tape migration is done. Available for local sources, disk sources, and standard FTP sources. This feature is not available for complex Virtual Objects.

- **Options**

  Additional options for performing the transfer of data from the Source Server to DIVA. These options supersede any options specified in the DIVA configuration database. Currently the possible values for Options are as follows:

  - No Entry

    No entry in this field specifies no options.

  - `-r`

    Using -r specifies that every name in *filenamesList* that refers to a folder must be scanned recursively. This also applies when *FilesPathRoot* is specified and an asterisk designates the files to be archived. This option can be used when archiving from a local source or from a standard FTP Server.

  - `-login`

    A user name and password is required to log in to some sources. This option obsoletes the -gateway option from earlier versions.

  - `-pass`

    The password used with -login.

- **Comments**

  Optional information describing the Virtual Object. This field is optional and can be left empty.

- **Additional Services**

  Select this check box to delete the original file after it has been archived.

**Note:** Delete on Source is not supported for Broadcast Servers.

- **Add New Component**

  Enter any additional components to be included in the Archive job.

- **Priority**

  The priority level selection for this job is in the form of a slider control, or text box at the bottom of the form page. The level can be in the range 0 to 100, or the value Default. If the Default Priority check box is selected, the slide control becomes inactive and the priority defined in the DIVA configuration is used.

  The value 0 is the lowest priority and 100 the highest priority. Move the slide control to increase or decrease the job priority, or enter the desired priority number in the Priority Value text box.

  There are six predefined values as follows:

  - `Min`
  - `Low`
  - `Normal`
  - `High`
  - `Max`
  - `Default`

  If the Default Priority check box is selected, the slide control becomes inactive and the priority defined in the DIVA configuration is used.

## Archive Job Files Path Root and Files Parameters

The Files Path Root and Files parameters in the Archive Submission form determines the main folder location, and the subfolders and files to be archived. Each serves a different purpose, yet both parameters work with each other. Identify a logical business object before filling in these parameters and executing the job.

The Files Path Root field identifies the path to the main file folder (the top folder). For example, C:\DROPFOLDER\Media\Virtual Object1\.

What you enter in the Files field text box identifies the individual files under the main folder (the identified Files Path Root) and any additional subfolders and files. For example, subfolder1\file3.

The Files field may contain an absolute path. However, this is not recommended because it prevents the object from being restored to a different root folder.

Assuming a Files Path Root is identified, do not use the full file path in the Files field. Only use the folder names and file names that are located under the identified Files Path Root folder. Alternatively, the Files Path Root field can be left blank and the full file path and name may be entered into the Files field.

The following are examples of how these parameters can be utilized:

## Correct Examples

The following entries archive only the specified files located in C:\DROPFOLDER\Media\Object1\ and the subfolder1\file3.

- `Files Path Root`

  *C:\DROPFOLDER\Media\Object1\*

- `Files`

  *file1*
  *file2*
  *subfolder1\file3*

The following entries archive all folders and files located in C:\DROPFOLDER\Media\Object1\.

- `Files Path Root`

  *C:\DROPFOLDER\Media\Object1\*

- `Files`

  *

The following entries are correct but not recommended, because in the future the object cannot be restored to a different location. The system loses flexibility and compatibility with other storage devices and in some scenarios, transcoding and Partial File Restore capabilities are also limited. In this example, the Files Path Root was left blank and the absolute paths are entered in the Files field.

- `Files Path Root`

- `Files`

  *C:\DROPFOLDER\Media\Object1\file1*
  *C:\DROPFOLDER\Media\Object1\file2*
  *C:\DROPFOLDER\Media\Object1\subfolder1\file3*

## Incorrect Example

The following entries result in an error and the Archive job does not complete:

- `Files Path Root`

  *C:\DROPFOLDER\Media\Object1\*

- `Files`

  *C:\DROPFOLDER\Media\Object1\file1*
  *C:\DROPFOLDER\Media\Object1\file2*
  *C:\DROPFOLDER\Media\Object1\subfolder1\file3*

## Archive Job with Delete on Source

There are instances where you must delete content, and possibly the parent folder, on a server. There are two options available to satisfy all possible scenarios:

- `-r`

  Recursive delete
- `-delete_fpr`

  Recursive deletion including the parent folder

The two options work either separately or together as indicated in the following workflow examples:

### Example 1

DIVA deletes the content of C:\source\root recursively because of these settings:

- `Files Path Root`

  *C:\source\root*
- `Files`

  *
- `Options`

  *-r*

### Example 2

DIVA deletes the content of C:\source\root recursively and the parent folder (root) because of these settings.

- `Files Path Root`

  *C:\source\root*
- `Files`

  *
- `Options`

  *-r -delete_fpr*

### Example 3

DIVA deletes only the content of C:\source\root because of these settings.

- `Files Path Root`

  *C:\source\root*
- `Files`

  *
- `Options`

## Example 4

DIVA deletes only the content of C:\source\root, and eventually the parent folder (root) if it is empty, because of these settings.

- `Files Path Root`

  *C:\source\root*

- `Files`

  *

- `Options`

  *-delete_fpr*

## Example 5

DIVA deletes the content of C:\source\root\Object recursively and the parent folder (Object) because of these settings.

- `Files Path Root`

  *C:\source\root*

- `Files`

  *Object\**

- `Options`

  *-r*

## Example 6

DIVA deletes the content of C:\source\root\Object recursively, then delete C:\source\root\Object, and finally delete C:\source\root if it is empty because of these settings.

- `Files Path Root`

  *C:\source\root*

- `Files`

  *Object\**

- `Options`

  *-r -delete_fpr*

## Example 7

DIVA deletes the content of C:\source\root\Object1 recursively, delete C:\source\root\Object1, delete the content of C:\source\root\Object2 recursively, and delete C:\source\root\Object2 because of these settings.

- `Files Path Root`

  *C:\source\root*

- `Files`

  *Object1**
  *Object2**

- `Options`

  *-r*

### Example 8

DIVA deletes the content of C:\source\root\Object1 recursively, delete C:\source\root\Object1, delete the content of C:\source\root\Object2 recursively, delete C:\source\root\Object2, and delete C:\source\root if it is empty because of these settings.

- `Files Path Root`

  *C:\source\root*

- `Files`

  *Object1**
  *Object2**

- `Options`

  *-r -delete_fpr*

### Example 9

DIVA deletes the content of C:\source\root\Object1 recursively, delete C:\source\root\Object1, delete C:\source\root\Object2\subfolder\clip.mov, delete C:\source\root\Object2\subfolder if it is empty, delete C:\source\root\Object2 if it is empty, and delete C:\source\root if it is empty because of these settings.

- `Files Path Root`

  *C:\source\root*

- `Files`

  *`Object1\*`*
  *`Object2\*`*

- `Options`

  *-r -delete_fpr*

## Restore Jobs

A Restore is defined as the transfer of an object to a Target Server. Restore jobs can be initiated from the web app Content Management > Catalog Browsing page. Click the three dots next to the object to restore and select the restore type from the context menu. Fill out all required information on the form and click Submit to execute the job.

This process submits an object Restore job to DIVA and the DIVA chooses the appropriate instance to be restored. The job fails if the requested object is on media that is not available.

The following fields are included in the Restore Job form:

- **Object Name**

  The name of the object to be restored.
- **Object Collection**

  The Collection assigned to the object when it was archived. This parameter can be left empty but this may result in an error if several objects have the same name.
- **Instance**

  If multiple instances of an object reside in DIVA, the particular instance to restore can be specified. If left blank, DIVA selects the instance that provides the most optimum transfer.
- **Destination Servers (required)**

  Destination (for example, a video server or browsing server) for the object files. This name must be known by the DIVA configuration. Use the drop-down list to select the desired Destination.
- **Quality Of Service (required)**

  One of the following codes (see *Quality of Service* for detailed descriptions):

  - Default

    Restoring is performed according to the default Quality of Service (currently direct and cache for restore operations).
  - Cache Only

    Use cache restore only. Cache only restores can only restore from a tape instance. However, a tape instance on a tape in a Tape Group with a higher priority is preferred to a tape instance on a tape in a Tape Group with a lower priority.
  - Direct Only

    Use direct restore only—no disk instance is created.
  - Cache and Direct

    Use cache restore if available, or direct restore if cache restore is not available.
  - Direct and Cache

    Use direct restore if available, or cache restore if direct restore is not available.
  - Nearline Only

    Use Nearline restore only. Nearline restore restores from a disk instance if one exists. Otherwise, it creates a disk instance and restore from the newly created disk instance. However, a disk instance on a disk in an array with a higher priority is preferred to a disk instance on a disk in an array with a lower priority.
  - Nearline and Direct

Use Nearline restore if available, or direct restore if Nearline restore is not available. However, a disk instance on a disk in an array with a higher priority is preferred to a disk instance on a disk in an array with a lower priority.

- **Additional Services**

  Additional and optional services are available using the Additional Services pull-down list (to the right of the Quality of Services). To request those services select from the following options:

  - Default

    Operate using the default setting in the DIVA configuration.

  - Do Not Overwrite

    Do not overwrite existing files on the destination server.

  - Do Not Check Existence

    Do not check for the existence of the clip on the server.

  - Delete And Write

    Force delete and rewrite if Virtual Object exists on the server.

- **Options**

  Additional options for performing the transfer of data from DIVA to the Target Server. These options supersede any options specified in the DIVA configuration database. Currently the possible values for Options are as follows:

  - No Entry

    No entry in this field specifies no options.

  - `-login`

    A user name and password is required to log in to some sources. This option obsoletes the -gateway option from earlier versions.

  - `-pass`

    The password used with -login.

- **Files Path Root**

  The root folder on the destination where the object files are placed. This option appends or overrides the FPR used in the original archive job. If left empty, the files are placed in the Files Path Root folder specified when archiving the object.

- **Priority**

  The priority level for this job. The level can be in the range 0 to 100, or the value Default. The value 0 is the lowest priority and 100 the highest priority. Move the slide control to increase or decrease the job priority, or enter the desired priority number in the Priority Value text box. If the Default Priority check box is selected, the slide control becomes inactive and the priority defined in the DIVA configuration is used.

There are six predefined values as follows:

- `Min`
- `Low`
- `Normal`
- `High`
- `Max`
- `Default`

If the Default Priority check box is selected, the slide control becomes inactive and the priority defined in the DIVA configuration is used.

## Archiving and Restoring in AXF Mode

An archive job for an AXF file results in DIVA automatically detecting that the file is an AXF file. Instead of archiving the AXF file itself, DIVA archives the contents of the AXF file, retrieving the Checksums and Provenance of the object.

The restore job optional parameter `-axf` instructs DIVA to restore the original asset into an AXF file. Instead of purely restoring the content of an object to the destination, DIVA restores the content into a new AXF Wrapper. When combined with `-rm` or `-rxml`, this option can be used to export an object with metadata information and then drop it into a WFM Watch Folder.

The AXF archive and restore functionality includes the following:

- Archive the content of an AXF file using auto-detection.
  - Identifies the *axf* file name extension
  - Confirms it is a single file
  - Checks the beginning of the file for specific AXF properties
  - Checks for metadata information
- Restores an object into a new single AXF file. Previously this operation would have resulted in multiple files.
- Preservation of checksums
- Preservation of metadata
- Preservation of provenances
- Complex Object support

This options works with FTP_STANDARD, LOCAL, DISK, CIFS, and EXPEDAT Servers.

## Staging Restore Jobs

For content that is already available in the target media, DIVA does nothing.

For content that does not exist on the target media, DIVA performs a Copy to Tape Group to make an instance available on the target media. If the target media is an object storage, there is an optional parameter to specify the desired storage class (see *REST API Parameters*).

If the content is available in archive-type storage, such as Glacier, DIVA stages the content to the target storage class. The number of days in which the content is available in the target storage class following a stage job is be configurable (see *REST API Parameters*). The staging job may include a parameter to specify the number of days a restored object must be available. By default it is set to 1 day.

If the content is on a cloud bucket in archive-type storage such as Glacier, and the content has already been restored, DIVA checks the restoration expiry date and may eventually extend it if the number of days available requested is beyond the expiry date.

For validation purposes, the restoration expiry date can be verified on the S3 console to validate the functionality:

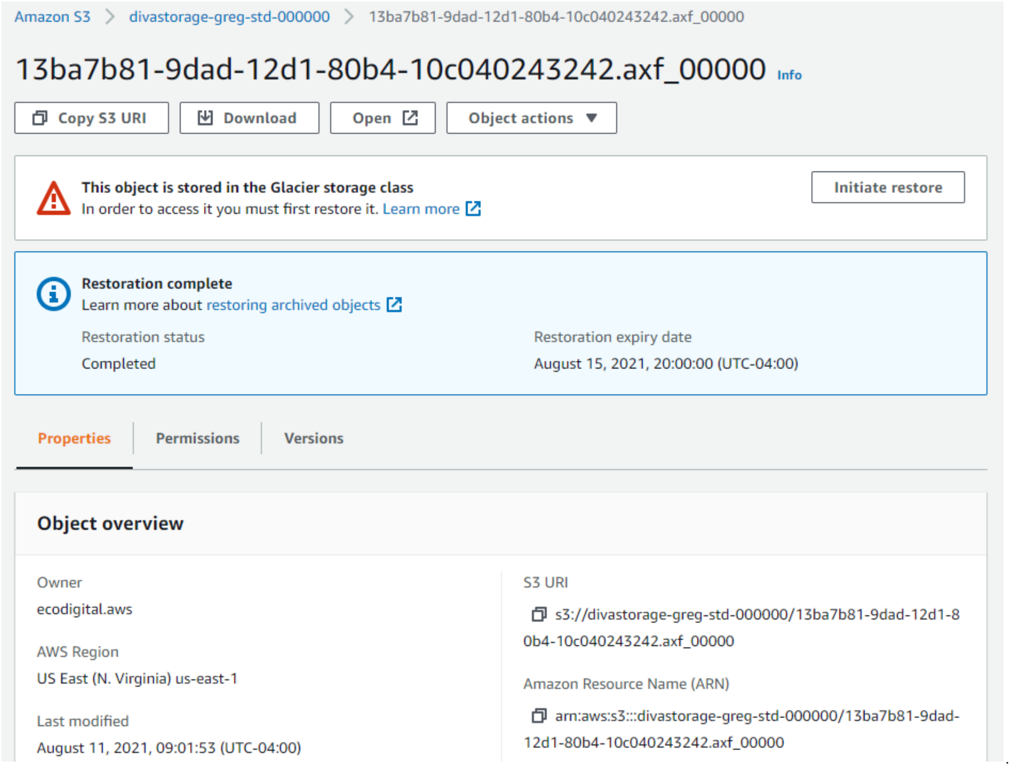
Amazon S3 > divastorage-greg-std-000000 > 13ba7b81-9dad-12d1-80b4-10c040243242.axf_00000
Info
Copy S3 URI
Download
Open
Object actions
Initiate restore
Properties
Permissions
Versions

13ba7b81-9dad-12d1-80b4-10c040243242.axf_00000

**This object is stored in the Glacier storage class**
In order to access it you must first restore it. Learn more

**Restoration complete**
Learn more about restoring archived objects

Restoration status
Completed

Restoration expiry date
August 15, 2021, 20:00:00 (UTC-04:00)

**Object overview**

Owner
ecodigital.aws

AWS Region
US East (N. Virginia) us-east-1

Last modified
August 11, 2021, 09:01:53 (UTC-04:00)

S3 URI
s3://divastorage-greg-std-000000/13ba7b81-9dad-12d1-80b4-10c040243242.axf_00000

Amazon Resource Name (ARN)
arn:aws:s3:::divastorage-greg-std-000000/13ba7b81-9dad-12d1-80b4-10c040243242.axf_00000

### REST API Parameters

A new endpoint supports the Stage job in the DIVA REST API. The POST endpoint accepts the name of the Object, Collection (or Collection Name), target media to stage the content, target storage class (if not specified it uses the default), restore tier, priority, options and number of days to stage the content. Here is a sample body:

```
{
  "Virtual ObjectName": "testVirtual ObjectName",
  "collectionName": "testCollectionName",
  "targetMedia": "testMedia",
  "numDaysAvailable": 3,
  "options": "",
  "priority": 50,
  "restoreTier": "EXPEDITED",
  "storageClass": "STANDARD"
}
```

### Failure Conditions

DIVA terminates a stage job under the following conditions:

- Object does not exist
- Target media does not exist
- Availability is beyond range.

DIVA configuration setting for regulating the upper limit for this value as follows:

- `#RELOADABLE in SERVICE mode`

  The number of days an instance is available after it is staged. The default setting of 3 implies a value between 1 day and 3 days is accepted.

- `#MAX_DAYS_FOR_STAGING=3`
  - Priority is in range (1-100)
  - Restricted character is not used
  - Object is available

## Partial File Restore Jobs

DIVA supports four types of Partial File Restore. The type implemented is determined by the job’s format parameter. This job submits a Partial Object Restore job and DIVA chooses the appropriate instance to be restored. The job fails if the requested object is on media that is not available.

The following list describes each type of Partial File Restore:

- **Files and Folders**

  This type of Partial File Restore enables extracting entire files from the archive or extracting entire directories and their contents. Multiple files and directories can be extracted in the same job. The files are restored with the file names and path names that were specified in the archive. There is no valid renaming option in File and

Folder Partial File Restore. For example, a file archived as misc/12-2012/movie.avi is partially restored to a misc/12-2012 subdirectory with the name movie.avi.

When a folder is specified in a File and Folder Partial File Restore, all files within that folder (and the folder itself) are restored. Also, each folder can include the `-r` option to recursively restore all folders nested within the target folder.

- **Byte Offset**

  This type extracts a range of bytes from a particular file in the archive. For example, you can extract bytes 1 to 2000 (the first 2000 bytes of the file), or byte 5000 to the end of the file (or both) and restore them to an output file such as movie.avi.

**Note:** The result of the Byte Offset Partial File Restore is usually not playable when applied to video files. The Actor does not apply the header, footer, etc., according to the video format.

- **Timecode**

  This type of Partial File Restore enables selecting a portion of a particular media file based on a timecode. For example, a user could extract from 00:00:04:00 to 00:10:04:00 (a 10 minute segment starting 4 seconds in and ending at 10 minutes and 4 seconds), and place that segment into an output file such as movie.avi. The resulting file is a smaller version of the original movie file.

**Note:** The result of the Timecode Partial File Restore is a valid clip when applied to video files. The Actor applies the header, footer, etc., according to the video format. If the Actor cannot parse the format, the job is terminated. This type of Partial File Restore can only be applied to a valid video clip.

- **DPX**

  This type of Partial File Restore enables extracting a range of DPX files from the archive. The entire object is viewed as a single media item, with one DPX file representing one frame of media. Only files with a *dpx*, *tif*, or *tiff* extension in the archive are considered frames for the purposes of this command.

  The first dpx | tif | tiff file in the archived object is considered frame 1, the second file in the archive is frame 2, and so on.

  In the unlikely event that these files are mixed, the first sequential file of any of the three extensions determines which files are considered to be part of the sequence. For example, if a stray tif file is mixed with a collection of dpx files and it came first in the sequence, the sequence is interpreted as a tif sequence and dpx files are ignored, even if this was not the intention.

  For example, to extract frames 10 through 15 using DPX Partial File Restore, it restores the tenth dpx file that appears in the archive, the eleventh dpx file, and so on, ending with the fifteenth dpx file, for a total of six files. Any other files (such as WAV files) are skipped by DPX Partial File Restore.

Special frame numbers 0 and -1 may be used to refer to the first and last frame respectively. Frame 0 is valid as the start of a frame range and Frame -1 is valid as the end of a range.

Valid frames and ranges are as follows:

- Frame 0 = first frame (select the Start of File check box)
- Frame 1 = the first frame in the sequence
- Frame n = the nth frame in the sequence
- Frame -1 = last frame (select the End of File check box)
- Specifying Frame 0 as the last frame is considered invalid.
- Specifying Frame 0 to 0 is invalid; won't return the first frame as may be intended.
- Specifying Frame 0 to 1 or Frame 1 to 1returns the first frame.
- Specifying the Frame -1 in the first frame currently produces an error. You also cannot specify Frame -1 to -1 to return the exact last frame if the exact number of the last frame is unknown.

**Examples**:

- `startRange=0 - endRange=1`

  Restores only the first frame.
- `startRange=600 - endRange=635, startRange=679 - endRange=779`

  Restores frames 600 through 635, and frames 679 through 779.
- `startRange=810 - endRange=-1`

  Restores all frames from frame 810 to the end of the archive.

  The actual file name may (or may not) match the frame number in DIVA. After restore DIVA interrogates the archive, finds the file order, and determines the Frame Number from the resulting file order found, it does not consider the filename. The first dpx, tif, or tiff file found is considered Frame 1.

  Be careful when archiving DPX files to ensure they can be partially restored properly—DPX Partial File Restore does not examine the filename or the DPX header information to determine which file is assigned to which frame. The assignment is based on the order in which the DPX files appear in the archive. By default this order is based on ordering established by the source and is typically alphanumeric. For example, NTFS DISK Servers order files and folders case insensitively as a general rule, except where diacritical marks, such as ', `, ^, and so on are applied.

  By default, when DIVA encounters a subfolder it recursively processes all of the children of that folder (including subfolders) before continuing with other files. If a folder appears in the alphanumeric folder listing it is archived recursively in the order it appears, but this can potentially create some issues. For example, if you want all of the subdirectories of a given directory processed first, followed by the files in the directory. Or, you might want all files processed first, then subdirectories.

DPX Partial File Restore uses an entire object as a single piece of media. If multiple reels or clips are in an archive, they can be stored in folders and partially restored using File and Folder Partial File Restore, but to DPX Partial File Restore they are viewed as a single clip. If this is a desired effect, ensure that the directories are sorted alphanumerically in the order the frames should be arranged.

DIVA does not perform any special audio handling for DPX media (other than what might be embedded in DPX files). DIVA can support transcoding of DPX media, but a transcoder may change the file names and/or file order of the DPX archive.

## Submitting a Partial File Restore Job

A Partial File Restore job is submitted by in the web app on the Content Management > Catalog Browsing page by clicking the three dots next to the object and selecting Partial Restore on the context menu. A Partial Restore form is displayed. Fill out the required information for the job and click Submit to execute the job.

Use the following procedure to navigate through the wizard:

1. The form opens with the selected Object Name and Object Collection already filled in. Select the type of Partial File Restore to perform using the menu.

   Each type of Partial File Restore has different options, except for Files and Folders Partial File Restore, which does not have any options.

2. Double-click the files from the left pane to copy them to the right pane to add them to the job.
3. Click Next to proceed.
4. You must include additional parameters for these Partial File Restore types. Click the edit icon (blue) next to the object name after it was moved it to the right pane. A form displays to enter the required additional parameters:

   – `Byte Offset`

     No offset is entered until you open the Options dialog box and manually enter them. Add the required offset parameters and click Add to include them in the job; then click OK to return to the job wizard.

   – `Timecode`

     The File Format (required) list is enabled after selecting the Timecode Partial File Restore. Select the proper file format from the list, or leave it set to Autodetect.

     Click the edit icon (blue) next to the object name after it was moved to the right pane. A dialog displays; enter the required additional parameters. Add the required Offset parameters and click Add to include them in the job; then click OK to return to the job wizard.

   – `DPX`

     Click the edit icon (blue) next to the object name after it was moved to the right pane. A dialog display, so that you can enter the required additional parameters.

Add the required parameters and click Add to include them in the job; then click OK to return to the job wizard.

**Note:** Partial File Restore jobs for AVI format files must include the same offset range (TCin, TCout) for all object components (for example, clip.avi, clip_1.wav, clip_2.wav).

The following list describes the parameters in the final Partial File Restore job screen:

- **Instance**

  If there are multiple instances of an object, DIVA selects the instance which allows the job to complete in the least amount of time (for example, a disk instance is selected over a tape instance). Specifying an instance number in this field overrides this behavior and target the specific identified instance.

- **Destination**

  Destination (for example, a video server or browsing server) for the object files. This name must be known by the DIVA configuration. Use the drop-down list to select the desired Destination.

- **Files Path Root**

  The root folder on the destination where the object files are placed. If left empty, the files are placed in the Files Path Root folder specified when archiving the object.

- **Options**

  Additional options for performing the transfer of data from DIVA to the Target Server. These options supersede any options specified in the DIVA configuration database. Currently the possible values for Options are as follows:

  - No Entry

    No entry in this field specifies no options.

  - `-login`

    A user name and password is required to log in to some Source Servers. This option obsoletes the -gateway option from earlier versions.

  - `-pass`

    The password used with -login.

- **Quality of Service**

  One of the following codes (see *Quality of Service* for detailed descriptions):

  - `Default`

    Restoring is performed according to the default Quality of Service (currently direct for restore operations).

  - `Cache Only`

    Use cache restore only.

  - `Direct Only`

    Use direct restore only; no disk instance is created.

  - `Cache and Direct`

    Use cache restore if available, or direct restore if cache restore is not available.

  - `Direct and Cache`

    Use direct restore if available, or cache restore if direct restore is not available.

- **Additional Services**

  Additional and optional services are available. To request those services use the menu to select the following options:

  - `Default`

    Use the default settings.

  - `Do Not Overwrite`

    Do not overwrite existing files on the target server.

  - `Do Not Check Existence`

    Do not check if the object exists on the target server.

  - `Delete and Write`

    Force delete of existing instances on the target server and write a new instance.

- **Priority**

  The priority level for this job. The level can be in the range 0 to 100, or the value Default. The value 0 is the lowest priority and 100 the highest priority. Move the slide control to increase or decrease the job priority, or enter the desired priority number in the Priority Value text box.

  There are six predefined values as follows:

  - `Min`
  - `Low`
  - `Normal`
  - `High`
  - `Max`
  - `Default`

  If the Default Priority check box is selected, the slide control becomes inactive and the priority defined in the DIVA configuration is used.

6. After selecting the Partial File Restore type and associated options for each object, click the Submit button to execute the job.

## Multiple Restore Jobs

If an object is required on multiple destinations simultaneously, the Multiple Restore job enables specification of all necessary destinations in one command and submits it as a single job (as opposed to multiple standard Restore jobs for each destination). This is also beneficial when the restore involves a tape instance because the tape is accessed once for the transfer rather than multiple read operations for single restore jobs of the same object. Up to five simultaneous destinations are currently supported.

If the object to be restored is part of a spanned tape set, it must be restored to cache before the transfer to all destinations. If the transfer to one of the destinations fails, the others still proceed (if possible) and the job status is Partially Aborted.

If multiple renaming rules are defined, DIVA processes the rule for each server independently.

Use the following procedure to execute a Multiple Restore job:

1. Navigate to the web app’s Content Management > Catalog Browsing page.
2. Click the three dots next to the object to be restored. A form is displayed where required parameters can be set.
3. Enter, or select, all required parameters on the form.
4. Use the Destination Servers pull-down list to select the first destination server, then click the plus icon (green +) next to the Files Path Root text box. This adds the selected server to the list in the text area below Files Path Root.
5. Continue adding servers in the manner described in step 4 until all servers are selected.
6. Click Submit to execute the job.

## Delete and Delete Instance Jobs

Use the Delete command to delete all instances of an object, or only a specific instance of the object from DIVA.

**Caution: You must use this command with caution; data loss could result.**

This command submits an Object Delete job to DIVA and deletes every instance of the object (unless otherwise specified).

The Instance field of the Delete job determines exactly what is deleted from DIVA. If this field is left empty, then all instances of that object is deleted. A specific number entered into this field only deletes the specified instance.

Use the following process to delete an object:

1. Navigate to Content Management > Catalog Browsing.
2. Click the three dots next to the object name on the page.
3. Select Delete Object from the menu.
4. Fill in the fields on the form that displays and select the check box near the bottom indicating that it is understood that the instance is completely removed from the system.
5. Click Submit to process the job.

When the job completes, the Num Instances column is automatically updated.

**Note:** Deletes and repacks do not clear WORM media because these are Write-Once media. The instances are deleted but the space is not recoverable.

The following fields are included in the Send Delete job screen:

- **Object Name**

  The name of the Object to be deleted. This is automatically populated with the Object Name selected.

- **Object Collection**

  The Collection assigned to the object when it was archived. This parameter can be a null string, but this may result in an error if several objects have the same name.

- **Instance**

  If multiple instances of the Virtual Object reside in DIVA, use the pull-down list to specify which specific instance to delete. If no entry is selected in this field, DIVA deletes all instances of that object.

- **Media**

  The media can be an existing Tape Group or disk array. The menu only displays those items already configured in the DIVA configuration.

- Options

  Additional options for deleting the object from DIVA. These options supersede any options specified in the DIVA configuration database.

- **Priority**

  The priority level for this job. The level can be in the range 0 to 100, or the value Default. The value 0 is the lowest priority and 100 the highest priority. Move the slide control to increase or decrease the job priority, or enter the desired priority number in the Priority Value text box.

  There are six predefined values as follows:

  – `Min`

  – `Low`

  – `Normal`

  – `High`

  – `Max-`

  – `Default`

  If the Default Priority check box is selected, the slide control becomes inactive and the priority defined in the DIVA configuration is used.

### Deleting Instances on Cloned Tapes

Instances on tapes linked to a clone cannot be deleted. Attempts to delete an instance or object with instances on a Source or Clone Tape result in the job terminating. The clone Storage Link must be removed to delete the instance.

## Copy Jobs

Use the Copy command to create an instance of an existing object in the same or another Tape Group or array. This is useful for creating a backup copy of the object on another media.

Submits a job for copying an archived object to a new object, with another name and/or Collection, to the Manager and the Manager chooses the appropriate instance as the source of the copy. All types of transfers (disk to disk, disk to tape, tape to disk, and tape to tape) are supported.

In the event the requested object is on media that is not available, the job fails.

When a copy job is issued with no instance specified and there are multiple instances of that object, DIVA selects the instance that executes the copy operation in the shortest possible time (for example, a disk instance is selected over a tape instance). If an instance number is selected in the Instance field of the form, the copy operation uses that specific instance only.

Use the following process to copy an object:

1. Navigate to the Content Management > Catalog Browsing page.
2. Click the three dots next to the object to be copied.

3. Select Copy from the menu.
4. Fill in the fields in the form.
5. Click Submit to execute the job.

The following fields are included in the Copy job form:

- **Object Name**

  The name of the source object. This is automatically populated with the Object Name selected.

- **Object Collection**

  The Collection of the source object.

- **Destination Media (required)**

  The Destination Media can be either a Tape Group or Disk Array. Use the pull-down list to select the Destination Media to use. If the new instance is being created in the same tape group or array, the job only succeeds if it can be copied to a separate tape or disk.

- **Instance**

  If multiple instances of the object reside in DIVA, use the menu to select which specific instance to copy. If no instance is specified, DIVA selects the instance that provides the most optimal execution time.

- **Options**

  Additional options for copying the object in DIVA. These options supersede any options specified in the DIVA configuration database.

- **Priority**

  The priority level for this job. The level can be in the range 0 to 100, or the value Default. The value 0 is the lowest priority and 100 the highest priority. Move the slide control to increase or decrease the job priority, or enter the desired priority number in the Priority Value text box.

  There are six predefined values as follows:

  - `Min`
  - `Low`
  - `Normal`
  - `High`
  - `Max`
  - `Default`

  If the Default Priority check box is selected, the slide control becomes inactive and the priority defined in the DIVA configuration is used.

## Copy As Jobs

When an object is archived to DIVA, it is uniquely identified by its Object Name and Object Collection. Neither the Name nor the Collection can be altered once it exists within the Core database. The Copy As command allows creation of a new object in

DIVA with a new Object Name and/or Collection, and then the original object can be deleted if desired or necessary (the deletion must be performed manually).

Use the following process to copy an object as another object:

1. Navigate to the Content Management > Catalog Browsing page.
2. Click the three dots next to the object to be copied.
3. Select Copy As from the menu.
4. Fill in the fields in the form.
5. Click Submit to execute the job.

For any object, the Object Name may not necessarily match that of the file name of the essence stored within it. If you use the Copy As command to create an object, it still is restored using the same name with which it was originally archived.

For example:

If a file named *xyz* is archived from a server, regardless of what Object Name was given to it in DIVA, it always restores to a destination as *xyz* regardless of its Object Name.

When a Copy As job is issued, the Instance field of the job is automatically left empty (an instance number can be selected from the pull-down list before the job is issued). If this field is left empty when the job is submitted and there are multiple instances of that object, DIVA selects the instance that completes the transfer in the shortest possible time by default (that is, a disk instance is selected over a tape instance). This depends on the QOS specified and the DIVA configuration.

The following fields are included in the Copy As Job screen:

- **Object Name**

  The name of the source object. This is automatically populated with the Object Name selected.

- **Object Collection**

  The Collection of the source object. This is automatically populated with the Object Collection selected

- **Destination Object Name (required)**

  The name of the destination Virtual Object.

- **Destination Object Collection (required)**

  The Collection of the destination object.

- **Destination Media**

  The Destination Media can be either a Tape Group or Disk Array. If the new instance is being created in the same tape group or array, the job only succeeds if it can be copied to a separate tape or disk. Use the pull-down list to select the Destination Media.

- **Destination Storage Plan**

  The storage plan to assign to the new object on the destination. Use the pull-down list to select the appropriate storage plan.

- **Instance**

  If multiple instances of the object reside in DIVA, use the pull-down list to select which specific instance to copy. If no instance is specified, DIVA selects the instance that provides the most optimal execution time (for example, a disk instance is selected over a tape instance).
- **Comments**

  Comments added here are added to the new object's properties.
- **Priority**

  The priority level for this job. The level can be in the range 0 to 100, or the value Default. The value 0 is the lowest priority and 100 the highest priority. Move the slide control to increase or decrease the job priority, or enter the desired priority number in the Priority Value text box.

  There are six predefined values as follows:

  – `Min`

  – `Low`

  – `Normal`

  – `High`

  – `Max`

  – `Default`

  If the Default Priority check box is selected, the slide control becomes inactive and the priority defined in the DIVA configuration is used.

# Eject Tape Jobs

The Eject Tape job ejects the selected tapes from the associated library. One or more tapes can be selected simultaneously using the check boxes next to the tape barcode. Use the following process to eject a tape or tapes:

1. Navigate to Resources Management > Tapes
2. Select the tape or tapes to eject using the check boxes next to each desired tape.
3. Click the three dots and select Eject Tape(s) from the menu, or click the Eject Tape(s) button at the bottom of the tapes list.
4. Fill in all fields in the form displayed.
5. Click Submit to execute the job.

The following fields are included in the Eject Tape(s) form:

- **Comments**

  Comments can be added when the tape is ejected. These may refer to its location or other information. Comments can be viewed later by examining that tape’s properties in the Resources Management >Tapes page (click the tape barcode to view the tape properties).

- **Priority**

  The priority level for this job. The level can be in the range 0 to 100, or the value Default. The value 0 is the lowest priority and 100 the highest priority. Move the slide control to increase or decrease the job priority, or enter the desired priority number in the Priority Value text box.

  There are six predefined values as follows:

  – `Min`

  – `Low`

  – `Normal`

  – `High`

  – `Max`

  – `Default`

  If the Default Priority check box is selected, the slide control becomes inactive and the priority defined in the DIVA configuration is used.

### Ejecting Cloned Tapes

Ejecting cloned tapes works in the same way as exporting them; the associated tape is ejected as well. To disable this behavior, remove the clone storage link. Then, if only the clone tape is ejected, the source Tape must be set to not writable.

## Insert Tape Jobs

This job enables inserting a tape into a library through its CAP. The tapes can only be entered in the CAP after this command is issued with some library configurations.

**Note:** Contact Telestream (see *Telestream Contact Information*) for instructions on bulk loading of tapes into a library.

Initiate an Insert Tape job using the Insert button located under the tapes list on the Resources Management > Tapes page.

The Sony PetaSite PSC software enables you to enter a tape in its CAP and manually place it within the PetaSite. In this case, DIVA is not informed of the action by the PSC and does not recognize the tape until the library is audited using the web app.

Use the following process to insert a tape into a library:

1. Navigate to Resources Management > Tapes.
2. Click Insert at the bottom of the tapes list.
3. Fill in all fields in the form.
4. Click Submit to execute the job.

The following fields are included in the Insert Tape form:

- **Robot Manager Name**

  This list specifies the Robot Manager controlling the associated library for insertion of the tapes. Use the pull-down list to select the appropriate robot manager.

- **CAP ID**

  This list is for Managed Storage with multiple CAPs. Some Managed Storage does not unlock the CAP, enabling the tape to be inserted, until the Insert Tape command is issued. You can specify which CAP to unlock from this pull-down list.

- **Priority**

  The priority level for this job. The level can be in the range 0 to 100, or the value Default. The value 0 is the lowest priority and 100 the highest priority. Move the slide control to increase or decrease the job priority, or enter the desired priority number in the Priority Value text box.

  There are six predefined values as follows:

  - `Min`
  - `Low`
  - `Normal`
  - `High`
  - `Max`
  - `Default`

  If the Default Priority check box is selected, the slide control becomes inactive and the priority defined in the DIVA configuration is used.

## Repack Tape Jobs

The Repack Tape job sends a repack job for the selected or specified tape. A tape repack operation reclaims unusable space on a tape due to object deletions, and removes fragmentation.

**Caution: The Repack Tape function is not intended to move material from a tape that is already known to be generating read errors. Contact Telestream (see *Telestream Contact Information*) for advice in these situations.**

Use the following process to repack a tape:

1. Navigate to the Resources Management > Tapes page.
2. Click the three dots next to the barcode of the tape to be repacked, or select the check box next to the tape(s) to be repacked.
3. Select Repack Tape from the menu, or click the Repack Tape button at the bottom of the tapes list (if the check box method was selected).
4. Confirm the barcode on the form and select the job priority.
5. Click Submit to execute the job.

Tape repacking can be a lengthy process and DIVA, by default, considers tape repacks a low priority operation. If higher priority jobs are issued, a tape repack job can be temporarily suspended while the higher priority jobs are completed. If higher priority jobs are issued sporadically it can result in frequent mount and dismount operations of the drive performing the repack. Therefore, Telestream recommends that tape repack operations should be run during off-peak periods where the frequency of higher priority jobs is limited. Some installations may have a drive dedicated solely to tape repacking to prevent this scenario from occurring.

The repack process involves the following tasks (in order):

1. Mounting the source tape and reading all objects to temporary disk cache of an Actor enabled for repack operations.

**Note:** If the temporary disk cache is filled before reading all objects from the source tape, DIVA proceeds to Step 2 until the cache is cleared. DIVA then proceeds to read the remaining objects from the source tape. This process is repeated until all objects are read.

2. Mounting a tape from the Unused Tapes Sets pool associated with the Set ID of the Tape Group from the source tape.
3. Writing all objects from the temporary cache in Step 1.
4. Deleting the objects from the temporary cache after all objects have been successfully written to the new tape.
5. The original source tape is released to the Unused Tapes Sets pool and unassigned from the Tape Group.

If a read error occurs at some point during the repack process from the source tape or a write error occurs on the destination tape, the entire repack job is terminated and no objects from the source tape are deleted. If the cache filled during the repack job and objects were successfully written to another tape before the cache was cleared, those objects remain on the destination tape.

If a read error occurred, the source tape's repack status and write status is disabled. If a write error occurred, the destination tape has its write status disabled and won't be used for any tape write operations. The write and repack states of both tapes is viewable on the Content Management > Job History page.

During the manual repack of WORM media, the usual dialogs display, but a warning is included notifying you that the space on the source media isn't recoverable after the repack is complete. Deletes and Repacks do not clear WORM media because these are Write-Once media. The instances are deleted, but the space is not recoverable.

Use the same process previously described to also repack WORM media. The following fields are included in the Repack Tape form:

- **Repack WORM (Write-Once media) with barcode**

  Barcode of the tape to be repacked.

- **Priority**

  The priority level for this job. The level can be in the range 0 to 100, or the value Default. The value 0 is the lowest priority and 100 the highest priority. Move the slide control to increase or decrease the job priority, or enter the desired priority number in the Priority Value text box.

There are six predefined values as follows:

- `Min`
- `Low`
- `Normal`
- `High`
- `Max`
- `Default`

If the Default Priority check box is selected, the slide control becomes inactive and the priority defined in the DIVA configuration is used.

### Repacking Cloned Tapes

Tapes linked to a clone cannot be repacked. If a repack is attempted on a Source or Clone Tape, the job terminates. The clone Storage Link must be removed to repack either tape.

### Automatic Tape Repack

The Automatic Repack option allows configuring scheduled tape repack operations. Use the following process to enable automatic tape repacking:

1. Navigate to the Resources Management > Tapes page.
2. Select a tape, or tapes, using the check box next to the tape barcode.
3. Click the Automatic Repack button at the bottom of the tape list.
4. Select the Automatic Repack Enabled check box and fill in the information on the form.
5. Click Submit to enable automatic repacking of the selected tape(s).

## Verify Tape Jobs

The Verify Tape job initiates a system read-back through every object on the selected tape one at a time and verifies all of the checksum values.

Use the following process to verify a tape or tapes:

1. Navigate to the Resources Management > Tapes page.
2. Select the check box next to the tape(s) to be verified.
3. Fill in the required information in the form.
4. Click Submit to execute the job.

The following fields are included in the Verify Tape form:

- **Verify Tape with barcode**

  Barcode of the tape to be verified.

- **Priority**

  The priority level for this job. The level can be in the range 0 to 100, or the value Default. The value 0 is the lowest priority and 100 the highest priority. Move the slide control to increase or decrease the job priority, or enter the desired priority number in the Priority Value text box.

  There are six predefined values as follows:

  - `Min`
  - `Low`
  - `Normal`
  - `High`
  - `Max`
  - `Default`

  If the Default Priority check box is selected, the slide control becomes inactive and the priority defined in the DIVA configuration is used.

## Export and Import Tape Jobs

The Export Tape job type enables one or more tapes containing objects to be exported to another independent DIVA platform (for example, at a remote disaster recovery or partner site).

**Note:** See *Removable Media* for details on tape exporting and importing, bulk tape exporting and importing, encrypted tape exporting and importing, and API exporting and importing functionality.

The metadata of each tape (that is, the object names and collections it contains and their location on the tape itself) are maintained in the Core Database. For complex objects this information is also in the Metadata Database. Additionally, the metadata of each tape is saved to an XML file when the tapes are exported. The XML file transfers the metadata of each tape into the other DIVA platform's database when the tapes are imported.

The Export Tapes command is not used for transferring tapes between two or more Managed Storage controlled by the same Manager. Tapes (and the instances they contain) exported from DIVA using this command are also removed from the DIVA database. If the object being exported is the last (or only) instance of that object, it is removed entirely from the database.

The following parameters have been added for export and import tape functions. XML metadata files exported from all previous DIVA versions are not supported.

During an export and import of WORM media, whether the media is Write-Once and the media is a cartridge is identified in the exported XML file. This information is imported with the attributes *isWriteOnce* and *isCartridge* being either true or false.

Importing WORM media is supported by DIVA. When importing DIVA WORM media into a DIVA version earlier than DIVA Core 7.4, DIVA Core ignores the WORM flag (it is set to false), and logs it in the Manager log. The device is visible in the web app as a tape but unusable if finalized, or if no WORM drive is connected to the system.

The following table describes the export and import parameters:

| Parameter | XML Element or Attribute | Notes |
|---|---|---|
| ObjectId | Attribute of object element | Not imported. A new Object ID is generated during import. |
| uuid | Attribute of object element | Imported if present, otherwise a new UUID is generated. |
| numFolders | Attribute of object element | |
| format | Attribute of object element and attribute of the tape element | 0 = legacy<br>1 = AXF<br>-1 = unknown |
| numFolders | Attribute of object element | |
| isHeaderValid | Attribute of object element | |
| isComplex | Attribute of object element | |
| footerBeginPos | Attribute of the element | If exists in database |
| footerEndPos | Attribute of the element | If exists in database |
| compOrderNumBegin | Attribute of the element | If exists in database |
| compOrderNumEnd | Attribute of the element | If exists in database |
| fileFolderMetadataInfo | Element | Valid for complex objects |
| fileFolderMetadataInfo-elem | Element | Valid for complex objects |
| checksums and checksum | Element | Not valid for complex objects |

Use the following process to export tape(s):

1. Navigate to the Resources Management > Tapes page.
2. Click the three dots or select the check box next to the tape barcode.
3. Select Export Tape or click the Export Tape(s) button at the bottom of the tape list.
4. Fill in the required information in the in the form.
5. Click Submit to execute the job.

The following fields are included in the Eject Tape Job screen:

- **Comments**

  Enter any desired comments in this field.
- **Delete from DB**

  When selected removes selected tape(s) from the Exported Tapes list.
- **Exported Tapes**

  This area displays tapes selected for export. These tapes and the instances they contain are removed from the DIVA Database after being exported. If required, select any tapes to be removed from the Exported Tapes list and click Remove Selected.
- **Priority**

  The priority level for this job. The level can be in the range 0 to 100, or the value Default. The value 0 is the lowest priority and 100 the highest priority. Move the slide control to increase or decrease the job priority, or enter the desired priority number in the Priority Value text box.

  There are six predefined values as follows:

  - `Min`
  - `Low`
  - `Normal`
  - `High`
  - `Max`
  - `Default`

  If the Default Priority check box is selected, the slide control becomes inactive and the priority defined in the DIVA configuration is used.

## Exported Tape Metadata Files

DIVA writes each tape's metadata to an XML file when the tapes are exported from the system. If an object is spanned across two (or more) tapes, the XML file encompasses every tape in the spanned set. The naming format of each tape metadata XML file is `Tapeset-<Barcode>.xml` (for example, Tapeset-000131.xml).

The root path where the XML files are saved is defined by the DIVA_EXPORT_ROOT_DIR parameter in the DIVA configuration file (consult the System Administrator or Telestream Technical Support for these details). By default, the export absolute directory root path is %DIVA_HOME%\Program\DIVA Core\bin\exported\. From this root path, the XML files from each Export Tapes command are saved in subfolders based on the date and time the command was run.

## Exporting Cloned Tapes

Exporting a Source or Clone Tape also triggers the export of the associated tape. In a previous example, content was archived to Source Tape 3L2042 and the contents were cloned to tape 3L2048. When exporting Source Tape 3L2042, Clone Tape 3L2048 is automatically included in the list of tapes to export. Similarly, exporting Clone Tape 3L2048 triggers the export of tape 3L2042.

The resulting exported XML file(*s*) contains a new element named `cloneBarcode`. It is used during import to restore the clone Storage Link. It also has the `lastCloneDate` for the Source Tape that is also restored during import.

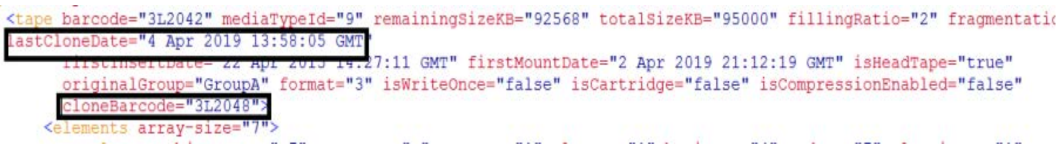

```
<tape barcode="3L2042" mediaTypeId="9" remainingSizeKB="92568" totalSizeKB="95000" fillingRatio="2" fragmentatic
lastCloneDate="4 Apr 2019 13:58:05 GMT"
      firstInsertDate="22 Apr 2019 14:27:11 GMT" firstMountDate="2 Apr 2019 21:12:19 GMT" isHeadTape="true"
      originalGroup="GroupA" format="3" isWriteOnce="false" isCartridge="false" isCompressionEnabled="false"
      cloneBarcode="3L2048">
    <elements array-size="7">
```

## Tape Import Workflow

To use the `importtapes` command, first ensure that the XML metadata file and the FFM files that were exported exist on the DIVA system into which they are to be imported. The files must exist in uncompressed form (unzipped) in DIVA's bin directory (by default). The Object Tape Group must also already exist on the target system before the import begins. This Tape Group does not necessarily have to be the same Tape Group that the tape was assigned to in the source system.

There are three main ways that a tape object can be treated during the import process:

- Imported as a new object
- Skipped if the object is already present in the system
- Imported as an instance of a preexisting (in the Core database) object. This option only functions correctly if the checksums match.

See the DIVA Installation and Configuration Guide for more information.

## Importing Tapes

The associated XML files from exported tapes must be copied to the DIVA bin directory on the platform where they are to be imported. Before inserting the tapes into the library, the tape metadata must be imported into the Core database. This must be executed for each tape (or spanned set) to be imported.

1. Open a command prompt window.
2. Execute the following commands in the order shown:

```
cd \%DIVA_HOME%\Program\diva\bin
importtapes <destination_group> <metadata_file>
```

– `destination_group`

The Tape Group that the tape (or spanned set) and its instances are assigned to in the destination system.

– `metadata_file`

The file name of the XML file for the tape (or spanned set).

### Example:

The tape with barcode number 000131 also contains objects that are spanned across the tape with a barcode of 000120. When tape 000131 is exported, the exported XML file is named Tapeset-000131.xml. This XML file also encompass the objects from tape 000120, and both tapes are ejected from the library.

After all objects from both tapes are exported to the XML file, all instances on each tape (and references to the tapes themselves) are removed from the Core database. The XML file is then copied to the %DIVA_HOME%\Program\diva\bin folder of the destination DIVA system where it is to be imported.

Import the metadata for this tape into the Tape Group MOVIES using the following command:

```
importtapes MOVIES Tapeset-000131.xml
```

After the tape's metadata successfully imports into the database (check the web app Job History page), both tapes and their objects are considered externalized and can both be entered into the library with the Insert Tape command.

## Importing Cloned Tapes

To restore the unidirectional Storage Link with the same Source and Clone Tape designations, the clone Storage Link must be imported first. Otherwise, the relationship becomes reversed; that is, the Source Tape becomes the Clone Tape and vice-versa. This is not critical because these tapes are clones of each other.

For example:

```
ImportTapes “GroupB”
“D:\diva\diva_std_730_jff\Program\Core\bin\exported\2019-04-04--
10.33.03\tapeset-3L2048.xml”
```

The import completed successfully. (Code 0)

Importing the associated tape requires the `-addAsInstanceIfNameExists` option, because the same objects are imported from the associated tape.

For example:

```
ImportTapes -addAsInstanceIfNameExists “GroupA”
“D:\diva\diva_std_730_jff\Program\Core\bin\exported\2019-04-04-
10.33.03\tapeset-3L2042.xml”
```

```
The following objects to be imported have the same Object Name and
Collection as objects already in the DIVA database... 1. A19/A 2.
A18/A 3. A17/A 4. A16/A 5. A7/A 6. A21/A 7. A20/A W A R N I N G ***
The 7 Virtual Object names listed above exist on the tape to be
imported, but ALSO exist in the DIVA system you are importing TO. *
You have chosen to first COMPARE the component CHECKSUMS of the
objects that have naming conflicts. If that match succeeds, these
objects will be imported as instances of the objects they match to.
* If the checksums match and you import as an instance, you may
lose the Comments, UUID, Archive Date and/or archived path root
```

```
that were stored in the import metadata. * If they don't match, the
import will fail, and then will exit...

The import completed successfully. (Code 0)
```

## Migrate Content Jobs

The Migrate job enables tape content migration to another Tape Group or Disk Array. For example, if upgrading to a new media type in your library and want to move the content from the old legacy tapes to the new AXF format.

**Note:** DIVA includes an embedded migration service (Migration Tool). It is a separate internal (to DIVA) service which helps users to schedule and run jobs to migrate content between different media inside of a DIVA system. The web app or command line client can be used. Contact Telestream (see *Telestream Contact Information*) for more details.

A migration job must be used to change a tape format from Legacy to AXF; repacking a tape does not change the tape format. Repacking of existing Legacy format objects retains the format of the tape even if the Tape Group format was updated in the configuration from Legacy to AXF.

This job type is only available from the web app Tapes View. It also uses the SPM (Storage Policy Manager) option to perform the migration (SPM must be installed). Appropriate slots must be configured for the migration. The slots indicate to DIVA when the migration operations are to be performed

See the Storage Policy Manager User Guide (https://www.telestream.net/telestream-support/diva/support.htm) for more information.

A Migrate job performs these functions (in the order shown):

1. Mount the source tapes and issue Copy jobs to DIVA to copy from the source tape to the destination media (disk or tape). Any spanned objects require the object to be copied first into cache.
2. Delete the Source Server instance after the object has been successfully copied to the new media.
3. The source tapes are cleared of all objects and returned to the Unused Tapes Sets pool.

Use the following process to initiate a migration job:

1. Navigate tot he Migrations > Start New Migration Task page.
2. Select either Copy or Move for the Migration Type.

3. Follow through the wizard for the selected migration type by clicking the appropriate selection on each of the following pages:

   a. Migration Source Type
   b. Migration Source
   c. Migration Target
   d. Migration Settings

   The pages automatically change when a selection is made. Some pages may require that you click the Next button at the bottom of the page to proceed to the next step.

4. When the Migration Summary page is reached confirm all of the selections and click Submit to execute the job. If changes are required, click the migration step number to go back to the step and make the necessary changes.

# Metasources

The Metasource source type allows several currently defined DIVA Servers sharing the same online storage (or monitoring the same folder or FTP server for Watch Folder Monitors) to be combined and considered a single DIVA Server configuration. This unique (and optional) feature enables DIVA to provide automatic load balancing and fail-over capabilities in case of one or more of the individual servers going offline.

When jobs are issued to DIVA with a server using the Metasource source type, each additional Archive or Restore job uses the next server in the Metasource list. If the server selected by DIVA is offline or encounters an error, DIVA automatically attempts to use the next server in the Metasource list. If all servers fail to fulfill the transfer, the job terminates.

# Symbolic Links

**Note:** This version of DIVA is for Windows only; there is no Linux support at this time.

Archive and restore symbolic links can be used on Linux. Symbolic links are only supported for the AXF format. When using LEGACY format, DIVA reports an error if a symbolic link is discovered during the transfer.

Symbolic Links are only supported with an SFTP Server on Windows. The following options must be specified when configuring SFTP:

```
-login [login] -pass [password] -port 22 -socket_block_size 64
```

When restoring an object containing symbolic links to a destination server that does not support them, they are ignored and not created on the destination server.

Symbolic links created using the Windows operating system are not supported. Shortcuts created using the Windows operating system are not represented as symbolic links because they are treated as files. Only symbolic links created on UNIX platforms are archived and represented as symbolic links in DIVA.

In the web app, the file type displays in the Content Management > Catalog Browsing > Object Properties page for the selected object under the Components area. The possible types are File, Directory, and Symbolic Link.

# Storage Policy Management

**WARNING: Misconfiguration of SPM (Storage Policy Manager) may lead to unexpected and disastrous results! Minor changes can lead to catastrophic consequences. For example, the deletion of hundreds of thousands of instances on tape or database corruption. Without special training and familiarity with the product, you should contact Telestream (see *Telestream Contact Information*) before making any changes to SPM. Failure to do so may result in severe damage to the DIVA system or even permanent data loss.**

The Storage Policy Manager software component enables object life cycle (interacting with the Manager) management, and is typically installed on the same system as the Manager. For example, an archived object can reside on a specific medium the first day, and migrate (over time) to a different medium according to the policies and rules established by you. DIVA executes the object life cycle migration as a background activity by following the rules and policies defined in the corresponding Storage Plan.

SPM supports disk cleaning based on the object's archived date. Earlier versions of SPM’s disk cleaning feature only supported cleaning based on an object's last access time and object size.

# Checksum Support and Content Verification

The purpose of the Checksum Support and Content Verification program is to introduce additional levels of verification into the DIVA system. This feature introduces checksum generation and verification for each file managed by DIVA. The currently supported checksum algorithms in DIVA include MD2, MDC2, MD5, SHA, SHA-1, and RIPEMD160.

**Note:** Additional checksum verification is performed at the Oracle Storage Cloud level. See the Storage Cloud documentation for information.

The default and recommended checksum algorithm is MD5. Although the other algorithms are maintained for backward compatibility, only MD5 and SHA-1 are recommended for best results.

When an object contains multiple files a checksum is generated and later verified for each of the component elements. The following types of checksum sources are distinguishable:

- Genuine Checksum
- Archive Checksum

- Deferred Checksum

The TEXT Genuine Checksum mode enables DIVA to archive all files and subfolders in a specified folder while comparing their checksum values against known values stored in an external checksum file. Files that do not have a matching checksum in the external checksum file are archived with DIVA's calculated checksum and the external checksum file is not archived.

**Note:** The TEXT Genuine Checksum is a customer-specific implementation and only supports MD5. Unicode is not supported, and checksums must be in an md5 text file.

## Web App Settings

Use the following process to configure general system checksum settings:

1. Navigate to the Configuration > General Settings > Checksums page.
2. Enable the checksum features on the Manager using the slide button (on or off).
3. Use the pull-down list to select the Default Checksum Type.
4. Enter the number of retries following a failed checksum in the text box.
5. Select whether the Manager should select a different Drive Per Retry on Failed Checksum using the slide button (on or off).
6. Click Save to save the configuration.

## Archive Instructions

Use the following procedure to archive objects using checksum verification through the web app:

1. In the web app, navigate to Content Management > Archive Submission.
2. From the Source drop-down list, select the Server entry that was created in the configuration stage.
3. Enter the desired File Path Root in the File Path Root text field.
4. Enter the path to the location of the checksum files in the Files field and append your entry with a wildcard symbol (an asterisk).
5. Enter *-r* In the Options field.
6. Enter the remaining parameters in the job form and click Submit.

### Limitations

The following limitations apply when using checksum verification.

- DIVA cannot open or create files on a Windows file system if their absolute path exceeds 256 characters. The Root Path must be no more than a total of 256 characters.
- Only ASCII, non-UTF-8 encoded checksum files are supported.

- Each line in the checksum file must begin with an MD5 checksum, followed by 2 spaces, and then the File Path to the referenced file.

## Genuine Checksum using AXF Transfer

The AXF Genuine Checksum mode enables DIVA to archive all files and subfolders in a specified AXF file while comparing their checksum values against known values stored in the AXF file. This kind of workflow is typically combined with a Restore job with `-axf` in the Job Options field.

### Requirements

The AXF Virtual Object containing the files to be archived must contain checksum information for each file. The supplied checksum in the AXF Virtual Object must be the expected type as defined in the configuration.

### Archive Instructions

Use the following procedure to archive an object with Genuine Checksum through an AXF transfer:

1. In the DIVA web app, navigate to Content Management > Archive Submission.
2. Fill in the required information in the text boxes.
3. From the Source Server drop-down list, select the Server entry that was created in the Configuration procedure.
4. Enter the desired File Path Root.
5. Enter the remaining parameters in the archive submission form and click Submit.

### Limitations

The workflow described only works with AXF jobs generated by DIVA.

Verify Following Restore (VFR) is not compatible with the -axf option. VFR was designed to read back the restored content from a video server to verify it has not been corrupted. The -axf option does not create an actual restore, but rather an object export in an AXF Wrapper. These options are mutually exclusive and should not be part of the same workflow.

# Quality of Service

The QOS (Quality of Service) parameter defines how a file is transferred to and from a tape, from a source, or to a destination. The following Job Options map to their logical quality of service:

- `-qos_direct_only`
- `-qos_cache_only`
- `-qos_direct_and_cache`

- `-qos_cache_and_direct`
- `-qos_nearline_only`
- `-qos_nearline_and_direct`

Job Options take precedence over the normal Quality of Service specification. Also, the normal Quality of Service specification takes precedence over the Server Connect Options.

NEARLINE_ONLY and NEARLINE_AND_DIRECT QOS values are supported in the Server Connect Options. These options are only valid for a Restore job. DIVA ignores the setting and the usual default is applied if a Source or Destination Server with either setting is used in any other type of job. The QOS value is not case-sensitive and no longer must be specified at the beginning of the options.

For example a valid option is:

```
-login test -pass test qos=nearline_only
```

The options for QOS are defined as follows:

- `Default`

  The QOS specified in the source or destination configuration is used.

- `Cache Only (-qos_cache_only)`

  The data is first transferred entirely to cache storage from tape, and then transferred to the destination. Alternately, the data is first transferred entirely from the source to cache storage, and then written to tape. If no cache service is available, the job terminates.

- `Direct Only (-qos_direct_only)`

  The data is transferred immediately to the source as it is read from tape. Alternately, DIVA writes the data to tape immediately as it is transferred from a destination. If no direct transfer service is available, the job terminates.

- `Cache and Direct (-qos_cache_and_direct)`

  If cache transfer is not available, for example no Actor with cache storage is available, then a direct transfer is performed.

- `Direct and Cache (-qos_direct_and_cache)`

  If a direct transfer is not available, for example no Actor with direct transfer enabled is available, then a cache transfer is performed.

- `Nearline Only (-qos_nearline_only)`

  **This is only available for Restore and N-Restore jobs.** If a Nearline disk instance exists, the data is transferred from Nearline disk to the destination. Alternatively, the data is first transferred entirely to Nearline storage on disk from tape, and then transferred to the destination. If no Nearline service is available, the job terminates.

- `Nearline and Direct (-qos_nearline_and_direct)`

  **This is only available for Restore and N-Restore jobs.** If Nearline transfer is not available, for example no Actor with Nearline storage is available, then a direct transfer is performed.

If an object to be restored has both disk and tape instances, and Cache Only, or Cache and Direct QOS is used, DIVA may restore the tape instance as first priority over the disk instance. This behavior depends on the DIVA_CACHE_QOS_USE_DISK setting in the DIVA configuration. If set to true, DIVA restores the disk instance regardless of the QOS specified.

The Cache transfer method is particularly important for optimum use of DIVA resources when the transfer speeds between tape devices and the server vary considerably. For example, if the tape drive in the job can write data at 400 Mbps, but the source can only deliver the data at 100 Mbps, the tape drive never achieves its optimum transfer rate. Using the Cache QOS, the file can be completely transferred to cache first, and the drive can complete its write operation at its maximum speed. This method enables the drive to be used for other jobs in a shorter time compared to the same transfer using the Direct QOS.

If an object to be restored has a disk instance the Nearline Only or Nearline and Direct QOS restores from that instance. If an object to be restored has only tape instances, the Nearline Only or Nearline and Direct QOS attempts to create a permanent disk instance and then restore from that instance. Every subsequent Nearline restore for the same object is blocked and waits for the first restore process to create a disk instance. If the first restore fails to create a disk instance the process repeats with the next restore attempting to create a disk instance. All other restores are blocked until the disk instance has been created.

The default QOS for Restore and N-Restore jobs is Nearline and Direct. If the restore job is a Transcode Restore, or if the destination server is a Movie2Me server, the Manager switches the restore QOS to Direct Only. Other QOS types are not supported in this case.

# Components

This DIVA version is a Windows-only release; there is no Linux support at this time. Refer to the DIVA Archive, Concepts and Glossary book for details on DIVA Architecture.

## Topics

# Web App

**Note:** DIVA Command has been deprecated starting with DIVA Core 8.3 and is replaced with the DIVA web app to configure and operate a DIVA system. Telestream recommends using the Chrome browser to access the DIVA web app.

The web app is a browser-based application used to monitor, control, and supervise operations in DIVA. Several web apps can be running and connected to the same DIVA system at the same time.

Access the web app in a browser at https://<ip_address>:8765/DIVAWebUI/.

# Password Security

You cannot use the default password to log in to the web app with the administrator or operator profile after DIVA Core installation is complete. You must assign an administrator or operator password in the Configuration Utility before you are permitted to switch to the respective mode in the web app.

If you attempt to switch to administrator or operator mode in the System Management App without first assigning a password to the respective profile, a dialog box is displayed notifying you that you must set a password for the corresponding profile in the Configuration Utility. After you set the corresponding profile password in the Configuration Utility the first time, it no longer matters what you use for the old password when changing passwords.

# Manager

The Manager is the main component in a DIVA system. All archive operations are controlled and handled by the Manager. Operation jobs are sent by initiator applications through the Client API. As a purchasable option, Manager also supports Main and Backup systems.

The Manager runs as a Windows Service. The service can be managed through the Windows Services screen. The static configuration file for the DIVA is diva.conf. Most settings in this file can typically be left set to the default values. Operations of DIVA can be monitored by launching the web app.

The batch files in the DIVA's bin folder can be used to perform the following major operations:

- Start, stop, and restart the DIVA Service. All of these operations can be executed using the DIVA batch file by specifying start, stop, or restart after the diva.bat command respectively (for example, diva.bat start).

  You can also terminate all jobs with a *graceful_shutdown* command. The *graceful_shutdown* command waits until all jobs have terminated before stopping the Manager instead of the abrupt shutdown that occurs with the stop command.

- Notify the Manager of any changes to the DIVA's configuration using the Notify-Manager batch file.
- Import tapes into a Tape Group using the importtapes batch file.
- List all active connections and end some connections (by connection identifier) with the ConnMgr batch file.

The diva.bat file enables running DIVA as a service or using a console window. Execute the batch file using the following command and parameters:

```
%DIVA_HOME%\Program\diva\bin\diva.bat [command] [options]
```

For example:

```
%DIVA_HOME%\Program\diva\bin\diva.bat start -conf
config_file_name.conf
```

Appending the -conf (or -f) option after one of the following commands specifies a specific configuration file to load settings from. The diva.bat command parameters are as follows:

- `install` (`-i`)

  Installs DIVA as a system service.
- `uninstall` (`-u`)

  Removes DIVA service.
- `start`

  Starts DIVA.
- `stop`

  Stops DIVA immediately if it is running.
- `graceful_shutdown`

  Stops DIVA after all jobs running at the time of the shutdown have terminated and ignores any new jobs.
- `restart`

  Stops and subsequently starts DIVA.
- `reload`

  Requests that the current service reloads its settings.
- `status`

  Determines whether the service is running and displays the status.
- `dump`

  Requests that DIVA Service create a system dump.
- `version` (`-v`)

  Displays DIVA version information and then exits.
- `help` (`-h`)

  Displays help information and then exits.

Refer to the DIVA Installation and Configuration Guide for information about running Windows services.

# Actor

The Actor is the data mover between devices in the network. It supports the data transfer between many different types of devices and handles transcoding operations with Telestream transcoding software (optional). All Actor operations are initiated and coordinated by DIVA. One or more Actors can be configured to be controlled by a single DIVA.

Each Actor runs as a Windows service and automatically starts and begins accepting connections from DIVA when the Actor host is started. Actor services on each host may be managed from the Windows services dialog box.

When restoring the same file to the same destination twice in parallel, the behavior on Windows and Linux is different. On Windows, the first restore (they cannot arrive exactly at the same time) locks the file so that the second one can terminate. On Linux, there is no such lock at the file system level. Both restores are executed at the same time, and both write to the same file. The content of the resultant file is not predictable.

**Note:** Linux-based Actors currently only support Telestream Vantage transcoding operations.

The following list are the Actor executable files:

- `%DIVA_HOME%\Program\Actor\bin\ActorService.exe command [option]`

  Executes commands for the Actor Service. Appending the `-conf` (or `-f`) option after one of the following commands specifies a specific configuration file to load settings from. The ActorService.exe command parameters are as follows:

  – `install` (`-i`)

    Installs the Actor as a system service.

  – `uninstall` (`-u`)

    Removes the Actor service.

  – `debug` (`-d`)

    Starts the Actor in console mode.

  – `version` (`-v`)

  Displays the Actor version information and then exits.

  – `help` (`-h`)

    Displays help information and then exits.

- `%DIVA_HOME%\Program\Actor\bin\scandrive.exe`

  Identifies the tape drives in the system. There are no command-line parameters.

- `%DIVA_HOME%\Program\Actor\bin\TapeReadingUtility.exe`

  Opens the Tape Reading Utility, which enables manually reading the tape drives in the system. There are no command-line parameters.
- `%DIVA_HOME%\Program\Actor\bin\VideoAnalyser.exe`

  Opens the Video Analyzer Utility. This utility displays the internal structure of a video format by dropping video files to the appropriate top tab for that file type (for example, drop a MOV file on the MOV tab, an AVI file on the AVI tab, and so on). File information is displayed in the lower window panes. There are no command-line parameters.

# Robot Core

Although you can use DIVA to only manage disk storage, storage capacity can be further expanded by adding one or more tape Managed Storage. In these cases, the Robot Core module provides an intermediate software layer for DIVA to interact with many different types of tape Managed Storage. It is connected to DIVA through TCP/IP.

The Robot Core interfaces to the library using either a direct interface to the library itself (through native SCSI or SCSI over Fiber Channel), or through an intermediate Ethernet connection to the manufacturer's own library control software.

The Robot Core alerts DIVA when the collection of tapes in the associated library requires synchronization with the Core database. This functionality is specific to the SCSI Robot Core module, and detects potential tape inventory mismatches between the Core database and the library inventory.

Potential inventory issues are trapped by the Robot Core if the library sends unit attention codes 06h 00h 28h (inventory may be altered), or 06h 00h 29h (reset occurred). When this happens, the Robot Core notifies DIVA so that it resynchronizes the database with the content of the library.

The Robot Core Client is accessed using either the Robot Core Client (command line based) or the Robot Core Client GUI. The Robot Core Client GUI is a graphical interface making it easy to interact with the Robot Core.

**Note:** If intermediate robotics control software is installed between the Robot Core and the library under control (such as ACSLS, SDLC, or PSC), it must be running before starting the associated Robot Core.

When the Robot Core Client command-line interface is started it displays a screen similar to a Windows command line. The Robot Core Client is already started and only the commands necessary to perform the required operations, or to display the required information, need to be entered.

The following list are the Robot Core executable files:

**Caution: Although a Robot Core may be restarted while DIVA is running, an attempt to mount a tape to a drive while the Robot Core is offline may cause the drives to be set Out of Order.**

- %DIVA_HOME%\Program\RobotCore\bin\RobotCore.exe command [options]

  Executes commands for the Robot Core Service. Appending the `-conf` (or `-f`) option after one of the following commands specifies a specific configuration file to load settings from. The RobotCore.exe command parameters are as follows:

  - `install` (`-i`)

    Installs the Robot Core as a system service.

  - `uninstall` (`-u`)

    Removes the Robot Core service.

  - `debug` (`-d`)

    Starts the Robot Core in console mode.

  - `version` (`-v`)

    Displays the Robot Core version information and then exits.

  - `help` (`-h`)

    Displays help information and then exits.

- `%DIVA_HOME%\Program\RobotCore\bin\RobotCoreClient.bat [rmHost] [rmPort]`

  This is a command line utility to take control of the Robot Core if the DIVA system is down.

  - `rmHost`

    The remote host name for the connection.

  - `rmPort`

    The remote host port for the connection.

- `%DIVA_HOME%\Program\RobotCore\bin\RobotCoreGUI.bat`

  This is a GUI utility to take control of the Robot Core if the DIVA system is down.

# Checksum Support and Content Verification

The purpose of the Checksum Support and Content Verification program is to provide additional levels of verification for each file managed by a DIVA system.

During the archive process the checksum is generated automatically by the Actor and stored in the database. This checksum is not verified until an initial read-back or restore operation is performed.

Checksum verification occurs when transferring data from a server or when reading data from a Source Server or a storage medium. The latter occurs during the retrieval of a object from a storage medium during routine functions (Restore, Copy, Repack, Transcode, but not Partial File Restore), or during a read-back from storage (Verify-Following-Write feature), or from the Source Server (Verify-Following-Restore feature).

The checksum verifications and failures can be viewed through the web app by navigating to the Content Management > Catalog Browsing page and viewing the Checksum Status column.

**Note:** Additional checksum verification is performed at the Oracle Storage Cloud level. See the Oracle Storage Cloud documentation for information.

# Import Tapes Tool

The importtapes.bat batch file imports one or more tapes into a user-specified Tape Group in the DIVA system. The XML files created during the tape export must be specified as a command-line parameter.

This tool only imports the tape's metadata into the database and not the actual objects (or tape) themselves into the system. The tapes must still be inserted using the Insert Tape functions.

See *Removable Media* for more detailed information.

Execute the importtapes.bat batch file using the following command and parameters:

```
%DIVA_HOME%\Program\diva\bin\importtapes.bat [group_name]
[mfiledir] [mfiledir]
```

The importtapes.bat command parameters are as follows:

- *help (-h)*

  Displays help information and then exits.

- *group_name*

  The Tape Group to which imported tapes belong.

- *mfiledir*

  The XML files containing exported tape metadata or a folder that contains the files. The first mfiledir is required, additional entries are optional. Multiple files may be used as follows:

  mfiledir1 mfiledir2 mfiledir3 mfiledir4 (and so on)

- *-skipIfNameExists*

**Caution: This is an advanced option and not recommended for normal use! Use of this option causes the object on the tape to become invisible and DIVA uses only the visible object existing in the system.**

This is an advanced parameter and skips importing of objects with naming conflicts. Normally, if the object name exists the program stops and nothing is imported. This option enables skipping the existing object and continue to import the next object in the XML file.

- *-useImportDateAsArchiveDate*

  Forces use of the Import Date instead of the original Archive Date as the Archive Date of the object imported into DIVA.

# Client API

The Client API is a set of documented functions enabling external applications, acting as clients, to use the services offered by the DIVA system.

A library of client functions is provided and must be linked to each client application. These functions encapsulate client commands into DIVA job messages sent over a TCP/IP connection to DIVA.

The `getFilesAndFolders` API call is called successively to get the complete file and folder list. Normally the first time the method is called the `startIndex` is set to 1. For successive calls, set the `startIndex` to the `endIndex` as returned from the previous call. After all jobs have been returned, a call to this method returns an empty list.

Folders do not contain a checksum, but several checksums per file are provided if available including MD5, SHA1, and so on. The returned information identifies which of the checksums is the Genuine Checksum.

See the corresponding manuals for specifications and details on the use of each API. Various APIs are available as follows:

- REST API (recommended for enhanced functionality and required by the web app)
- C++ API
- Java API
- Enterprise Connect (Web Services API)
- DIVAprotectWS API

**Note:** The Web Services API is called DIVA Connect. See the DIVA Connect documentation (https://www.telestream.net/telestream-support/diva/support.htm) for information on installation, configuration, and operations.

# Auto-Discovery Agent

The Auto-Discovery GUI is hosted on the same server as the Publisher. To access the Auto-Discovery Agent, go to the root of the Publisher end-point to load the interface.

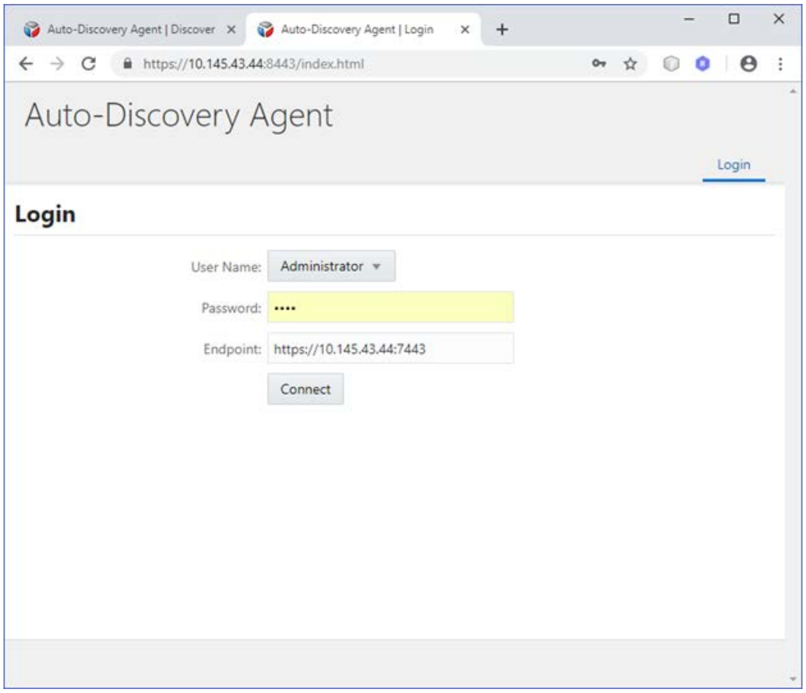
Auto-Discovery Agent | Discover
Auto-Discovery Agent | Login
https://10.145.43.44:8443/index.html
Auto-Discovery Agent
Login

Login

User Name: Administrator

Password: ••••

Endpoint: https://10.145.43.44:7443

Login as an Administrator or User from the main screen. A valid Data Service end-point must be specified to retrieve disk data from, and the Publisher end-point. After a successful connection the main status page is shown as seen here:

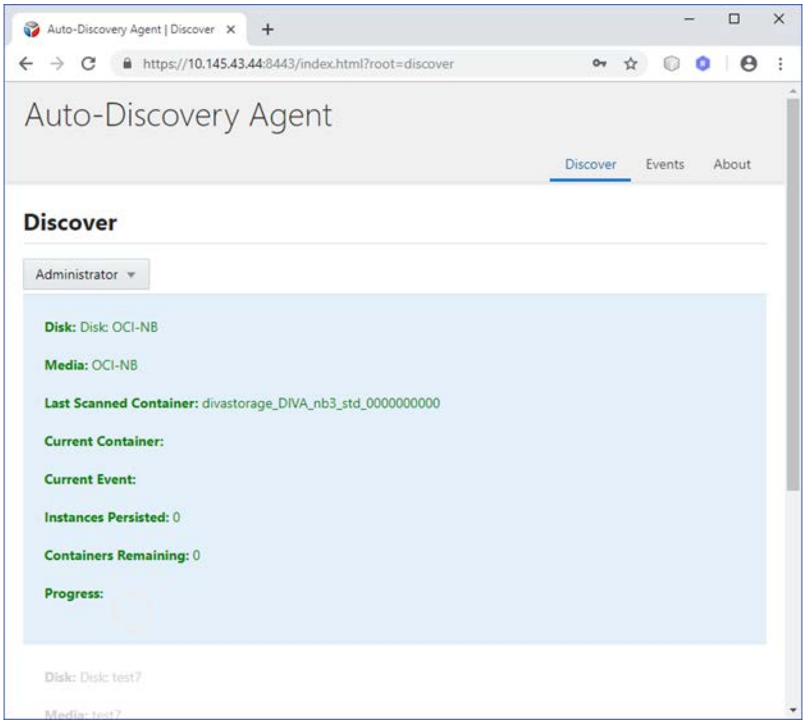
Auto-Discovery Agent | Discover
https://10.145.43.44:8443/index.html?root=discover
Auto-Discovery Agent
Discover
Events
About
Administrator

Discover

**Disk:** Disk: OCI-NB

**Media:** OCI-NB

**Last Scanned Container:** divastorage_DIVA_nb3_std_0000000000

**Current Container:**

**Current Event:**

**Instances Persisted:** 0

**Containers Remaining:** 0

**Progress:**

**Disk:** Disk: test7

**Media:** test7

Right-click on each disk and select one of the following operations from the pull-down menu: Start Scan, Resume Scan, Stop Scan or View Events:

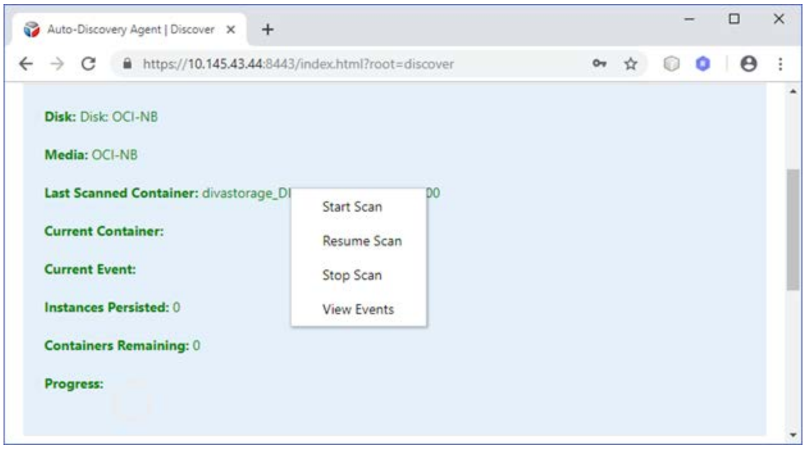
Auto-Discovery Agent | Discover
https://10.145.43.44:8443/index.html?root=discover
Start Scan
Resume Scan
Stop Scan
View Events

**Disk:** Disk: OCI-NB

**Media:** OCI-NB

**Last Scanned Container:** divastorage_DI 00

**Current Container:**

**Current Event:**

**Instances Persisted:** 0

**Containers Remaining:** 0

**Progress:**

Start Scan clears all persisted containers from the database and attempts to retrieve all content from the selected disk. Any instance that was previously persisted is skipped.

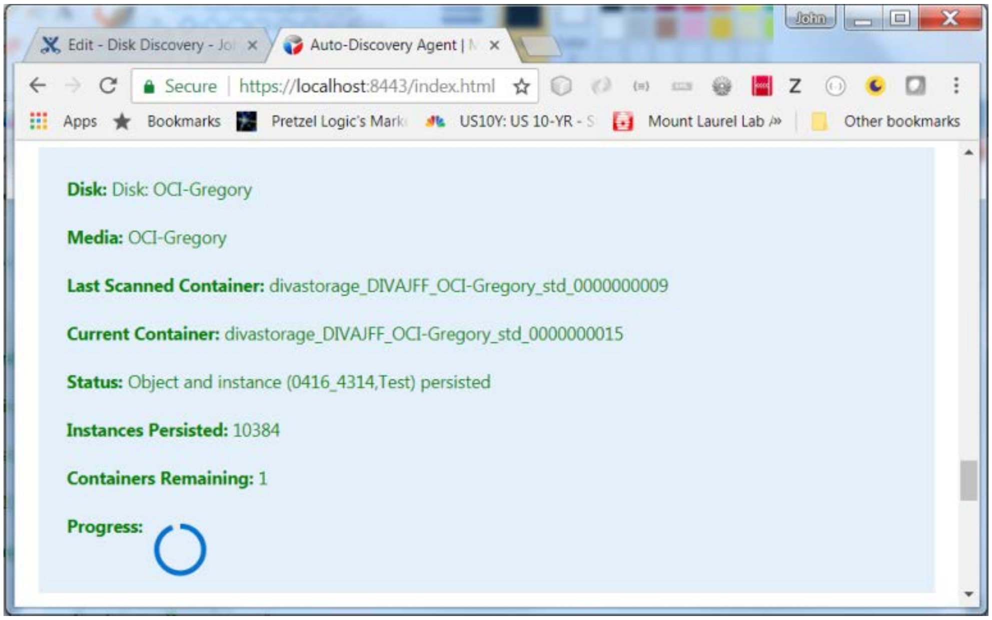

**Disk:** Disk: OCI-Gregory

**Media:** OCI-Gregory

**Last Scanned Container:** divastorage_DIVAJFF_OCI-Gregory_std_0000000009

**Current Container:** divastorage_DIVAJFF_OCI-Gregory_std_0000000015

**Status:** Object and instance (0416_4314,Test) persisted

**Instances Persisted:** 10384

**Containers Remaining:** 1

**Progress:**

Resume Scan attempts to resume a scan from the last successfully scanned container and instance.

Stop Scan instructs the Publisher to submit a stop request to all Actors currently scanning containers and terminate the disk scan. Unless a container was fully scanned, a Resume Scan operation attempts to continue a scan from the container on which a scan was terminated, starting from the discovered instance of that container with the largest UUID.

View Events displays the events screen shown in the following figure. This shows all of the events as stored in the database. The events are shown one page at a time:

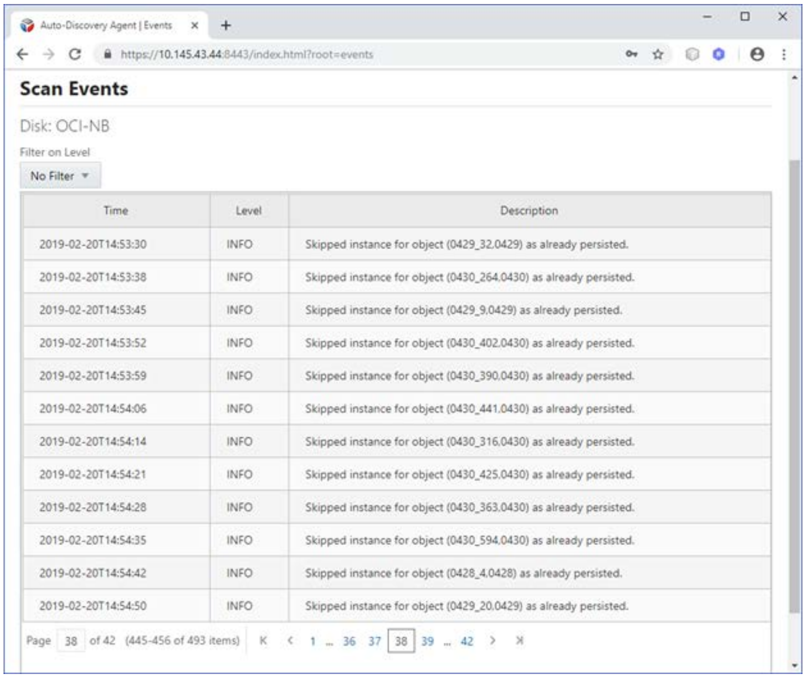
Auto-Discovery Agent | Events
https://10.145.43.44:8443/index.html?root=events

**Scan Events**

Disk: OCI-NB

Filter on Level

No Filter ▾

| Time | Level | Description |
|---|---|---|
| 2019-02-20T14:53:30 | INFO | Skipped instance for object (0429_32.0429) as already persisted. |
| 2019-02-20T14:53:38 | INFO | Skipped instance for object (0430_264.0430) as already persisted. |
| 2019-02-20T14:53:45 | INFO | Skipped instance for object (0429_9.0429) as already persisted. |
| 2019-02-20T14:53:52 | INFO | Skipped instance for object (0430_402.0430) as already persisted. |
| 2019-02-20T14:53:59 | INFO | Skipped instance for object (0430_390.0430) as already persisted. |
| 2019-02-20T14:54:06 | INFO | Skipped instance for object (0430_441.0430) as already persisted. |
| 2019-02-20T14:54:14 | INFO | Skipped instance for object (0430_316.0430) as already persisted. |
| 2019-02-20T14:54:21 | INFO | Skipped instance for object (0430_425.0430) as already persisted. |
| 2019-02-20T14:54:28 | INFO | Skipped instance for object (0430_363.0430) as already persisted. |
| 2019-02-20T14:54:35 | INFO | Skipped instance for object (0430_594.0430) as already persisted. |
| 2019-02-20T14:54:42 | INFO | Skipped instance for object (0428_4.0428) as already persisted. |
| 2019-02-20T14:54:50 | INFO | Skipped instance for object (0429_20.0429) as already persisted. |

Page 38 of 42 (445-456 of 493 items) K < 1 ... 36 37 38 39 ... 42 > ⟩|

The page is selected at the bottom of the table. You can use the provided filter to show only INFO, WARN, and ERROR events as displayed here:

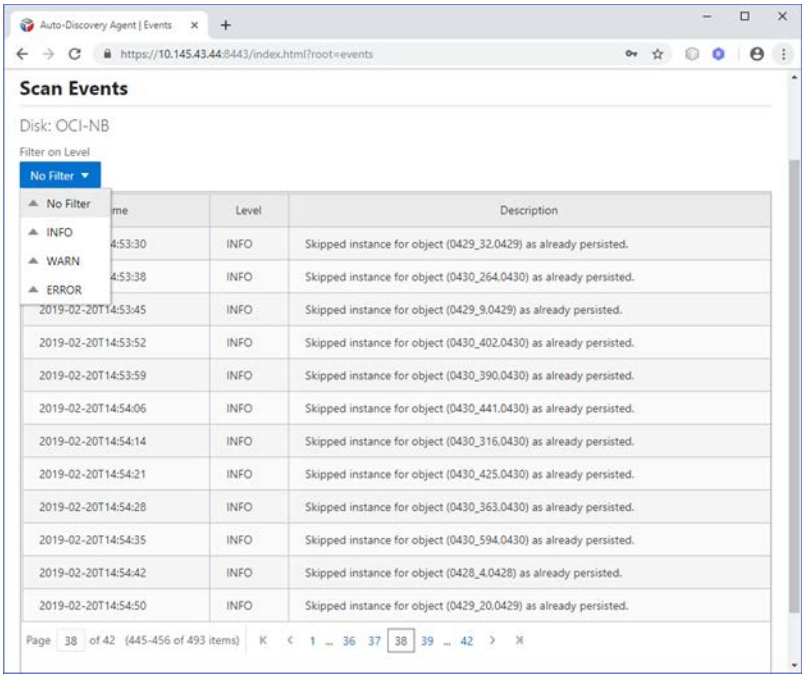
Auto-Discovery Agent | Events
https://10.145.43.44:8443/index.html?root=events
No Filter
INFO
WARN
ERROR

**Scan Events**

Disk: OCI-NB

Filter on Level

No Filter ▾

| me | Level | Description |
|---|---|---|
| 4:53:30 | INFO | Skipped instance for object (0429_32.0429) as already persisted. |
| 4:53:38 | INFO | Skipped instance for object (0430_264.0430) as already persisted. |
| 2019-02-20T14:53:45 | INFO | Skipped instance for object (0429_9.0429) as already persisted. |
| 2019-02-20T14:53:52 | INFO | Skipped instance for object (0430_402.0430) as already persisted. |
| 2019-02-20T14:53:59 | INFO | Skipped instance for object (0430_390.0430) as already persisted. |
| 2019-02-20T14:54:06 | INFO | Skipped instance for object (0430_441.0430) as already persisted. |
| 2019-02-20T14:54:14 | INFO | Skipped instance for object (0430_316.0430) as already persisted. |
| 2019-02-20T14:54:21 | INFO | Skipped instance for object (0430_425.0430) as already persisted. |
| 2019-02-20T14:54:28 | INFO | Skipped instance for object (0430_363.0430) as already persisted. |
| 2019-02-20T14:54:35 | INFO | Skipped instance for object (0430_594.0430) as already persisted. |
| 2019-02-20T14:54:42 | INFO | Skipped instance for object (0428_4.0428) as already persisted. |
| 2019-02-20T14:54:50 | INFO | Skipped instance for object (0429_20.0429) as already persisted. |

Page 38 of 42 (445-456 of 493 items) K < 1 ... 36 37 38 39 ... 42 > ⟩|

# Documented End-points

Both the Data Service and Publisher Service have a set of documented end-points where you can submit a job. All end-points of the Data Service except for `/swagger`, `/api-docs`, and `/auth-request` require an authentication token. POST a valid user name and password to the `/authRequest` end-point to obtain a token.

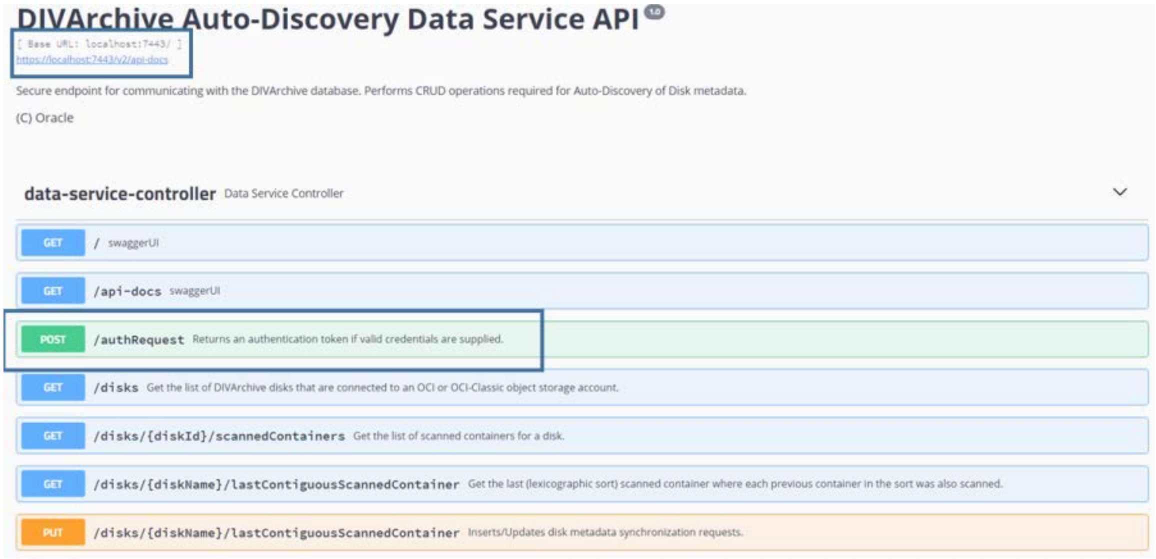

## DIVArchive Auto-Discovery Data Service API 1.0

[ Base URL: localhost:7443/ ]
https://localhost:7443/v2/api-docs

Secure endpoint for communicating with the DIVArchive database. Performs CRUD operations required for Auto-Discovery of Disk metadata.

(C) Oracle

**data-service-controller** Data Service Controller

- GET `/` swaggerUI
- GET `/api-docs` swaggerUI
- POST `/authRequest` Returns an authentication token if valid credentials are supplied.
- GET `/disks` Get the list of DIVArchive disks that are connected to an OCI or OCI-Classic object storage account.
- GET `/disks/{diskId}/scannedContainers` Get the list of scanned containers for a disk.
- GET `/disks/{diskName}/lastContiguousScannedContainer` Get the last (lexicographic sort) scanned container where each previous container in the sort was also scanned.
- PUT `/disks/{diskName}/lastContiguousScannedContainer` Inserts/Updates disk metadata synchronization requests.

You should copy and paste the response into the header of subsequent jobs:

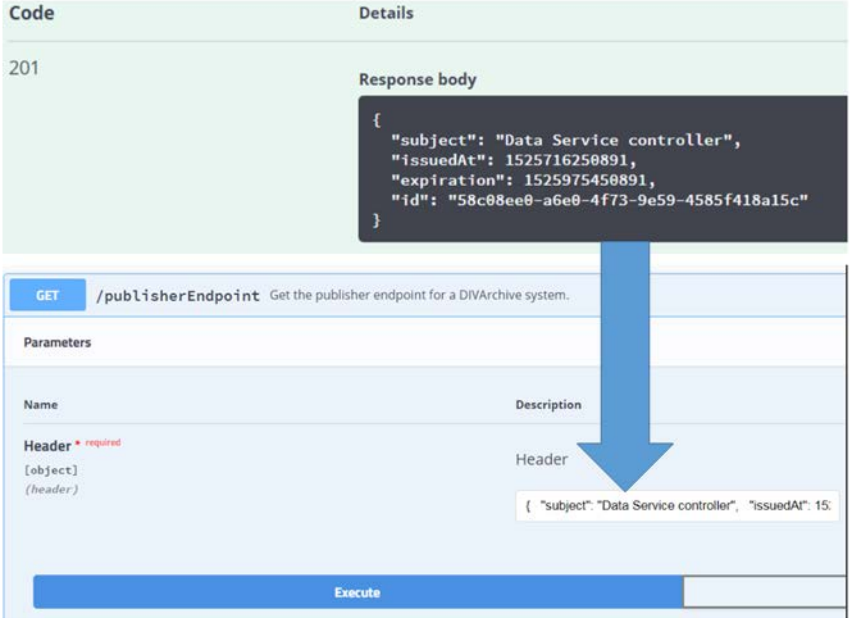

| Code | Details |
|---|---|
| 201 | Response body |

```
{
  "subject": "Data Service controller",
  "issuedAt": 1525716250891,
  "expiration": 1525975450891,
  "id": "58c08ee0-a6e0-4f73-9e59-4585f418a15c"
}
```

GET `/publisherEndpoint` Get the publisher endpoint for a DIVArchive system.

Parameters

| Name | Description |
|---|---|
| Header * required<br>[object]<br>*(header)* | Header<br>{ "subject": "Data Service controller", "issuedAt": 15: |

The Publisher end-point contains only two main end-points; one to submit a disk scan job (`/requests`) and the other to receive progress events (`/requests/progress`) as server-side events:

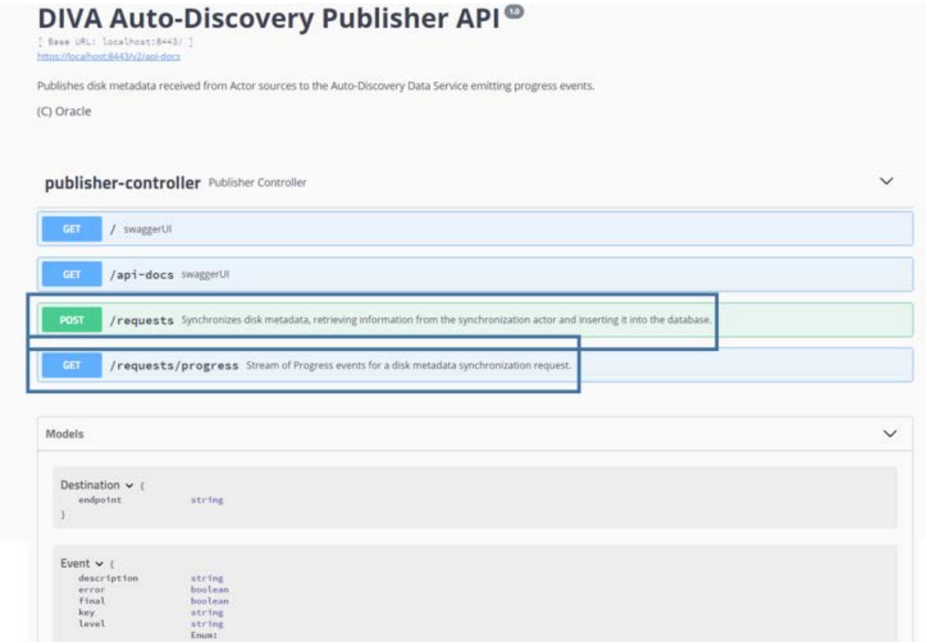

DIVA Auto-Discovery Publisher API 1.0

[ Base URL: localhost:8443/ ]

https://localhost:8443/v2/api-docs

Publishes disk metadata received from Actor sources to the Auto-Discovery Data Service emitting progress events.

(C) Oracle

publisher-controller Publisher Controller

GET / swaggerUI

GET /api-docs swaggerUI

POST /requests Synchronizes disk metadata, retrieving information from the synchronization actor and inserting it into the database.

GET /requests/progress Stream of Progress events for a disk metadata synchronization request.

Models

```
Destination {
   endpoint          string
}

Event {
   description       string
   error             boolean
   final             boolean
   key               string
   level             string
                     Enum:
```

# Watch Folder Monitor (Optional)

Watch Folder Monitor (WFM) provides automatic monitoring of newly-created files in up to 20 local or FTP folders (or combinations of the two). One (or multiple) files in FTP folders per object are supported. When a new file (or FTP folder) is identified, WFM issues an archive job automatically to DIVA to archive the new file or folders. After the files are successfully archived, they are automatically deleted from the Source Server.

When using WFM in a Linux environment to monitor an FTP folder, it must be configured as follows (this is an example):

- User: diva
- User home directory: /ifs
- Folder to be monitored: /ifs/folder1
- Correct WFM configuration: ftp://diva:password@host_ip/folder1
- Incorrect WFM configuration: ftp://diva:password@host_ip/ifs/folder1

The WFM Service can be started, stopped, and restarted using the operating system Services or the WFM command line utility on each host that is running a WFM installation. When the WFM Service starts, or restarts, WFM loads and validates the configuration file. If any configuration issues are detected, the process terminates and runs diagnostics.

If the configuration validation completes successfully, WFM begins scanning all of the configured Watch Folders, checks the status of all objects that were initialized before

WFM was last shutdown, and updates the internal database with the current status of the objects. After all of these checks have completed WFM is in the Running state.

When WFM finds files in a configured Watch Folder, it updates the internal database and requests DIVA to archive all files found as new objects. To avoid repeated archive jobs, WFM continuously updates the archive operations status in the internal database.

If jobs fail, the Status Module informs the internal database about the failure. If the number of unsuccessful job attempts reaches a pre-configured number, the object status is changed to "could not be archived" and the object is marked as incomplete. WFM logs information about the incomplete files and calls the WFM File DIVA Module to move them to the Trash folder.

If the job completes successfully, the internal database is updated by the Status Module. In the case of a File Set Object, WFM removes the Metadata File and the File Set folder.

WFM terminates upon execution of the shutdown script and stops all internal processes before all archive operations are complete. After all of the modules are stopped, all internal statuses on the disk are saved in the internal database before the WFM completes shutdown.

The WFM configuration file is %DIVA_HOME%\Program\conf\dfm\dfm.conf. Service logging is performed through the log file located in the %DIVA_HOME%\Program\log\dfm\ folder. The logging configuration is in the %DIVA%\Program\conf\dfm\dfm.trace file.

The dfm.bat file enables managing WFM from a command-line interface. Execute the batch file using the following command and parameters:

```
%DIVA_HOME%\Program\InterLink\dfm\bin\dfm.bat [command] [options]
```

Appending the `-conf` (or `-f`) option after one of the following commands specifies a specific configuration file to load settings from. The dfm.bat command parameters are as follows:

- `install` (`-i`)

  Installs the WFM module as a system service.

- `uninstall` (`-u`)

  Removes the WFM module service.

- `start`

  Starts the WFM module.

- `stop`

  Stops the WFM module if it is running.

- `restart`

  Stops and subsequently starts the WFM module.

- `status`

  Determines whether the service is running and displays the status.

- `version` (`-v`)

  Displays the version information and then exits.
- `help` (`-h`)

  Displays help information and then exits.

# SNMP Agent (Optional)

The SNMP (Simple Network Management Protocol) Agent and MIB (Management Information Base) supports status and activity monitoring of DIVA and its subsystems to a third party monitoring application through the SNMP protocol.

The SNMP agent is integrated with the Windows SNMP Service, which starts automatically when the server is started. SNMP information from DIVA to a monitoring application is obtained through the SNMP Agent, which in turn establishes a connection to DIVA automatically when DIVA is started.

**Note:** The SNMP Agent is not currently supported in the Linux environment.

Use the following procedure to configure the SNMP Service to monitor DIVA:

1. Install the Microsoft SNMP Services on the computer where DIVA is installed (if not already installed).

   **a.** On the server, navigate to the Windows Key > Administrative Tools > Server DIVA.

   **b.** Click Manage > Add Roles and Features.

   **c.** Click Next on each of the first four screens.

   **d.** Verify the SNMP Services are listed.

   If the SNMP Service is not listed in the services window, use the following procedure to add the service:

   **e.** On the server, navigate to the Windows Key > Administrative Tools > Server DIVA.

   **f.** Click Manage > Add Roles and Features.

   **g.** When you get to the screen showing services that can be installed, select SNMP Service and click Add Feature.

   **h.** Click Next > Install.

   **i.** After installation is complete, return to the Windows Services screen and click Refresh to refresh the display. The SNMP Service should now be included in the services list.
2. Stop the SNMP Service and SNMP Trap Service.
3. Navigate to the SNMP installation folder.

   ```
   %DIVA_HOME%\Program\SNMP\bin
   ```
4. Confirm this folder contains the DIVAapi.dll file. If not, you can copy it from the API Visual Studio .Net Dynamic Release directory.

5. Enter the correct DIVA connection information in the config.txt.ini file.

   Also, set the POLLING_RATE to 60 (to poll for jobs every 60 seconds), and remove the .ini from the end of the file name when saving the edited file.

6. Open the appropriate registry file and edit pathname so that it points to the SNMP path being used.

   For example, %DIVA_HOME%\Program\SNMP\bin\divasnmpagt.dll.

7. Double-click the registry file that was just edited to install the SNMP registry keys.
8. Start Regedit and using the registry information in the registry file, navigate to each registry key and make sure the path in the registry is as displayed in the registry file.
9. Open the SNMP Service properties and edit the following parameters:
   - On the Traps tab, enter *public* in the Community Name field and add the IP address of the computer where traps are viewed (for example, the computer where the MIB browser is installed).
   - On the Security tab, confirm the Send authentication trap and Accept SNMP packets from any host check boxes are selected.
   - In the Accepted community names field, add *public with READ ONLY rights*.
   - Click Apply.
10. Start the SNMP service. Do not start the SNMP Trap service.

The SNMP Service can also be manipulated through the Windows command prompt as follows (typically the same host as that of DIVA):

1. Open a Windows command prompt.
2. To start the SNMP Service, enter *net start "SNMP Service"* at the command prompt. The quotation marks are required for services with spaces in their service name.
3. To stop the SNMP Service, enter *net stop "SNMP Service"* at the command prompt. The quotation marks are required for services with spaces in their service name.

# Email Notifications

Email notifications, the notification frequency, and the time before the first email is sent are configurable in the DIVA web app on the Configuration > General Settings > SMTP Notifications page. An email is sent to the account name if any one of the following conditions occurs:

- A configurable minimum disk space constraint is reached
- A configurable minimum number of empty tapes is reached
- A maximum number of aborted jobs occurs
- The Actor-Disk connection goes offline
- The Actor-Drive connection goes offline
- A disk goes offline
- A drive goes offline

- An Actor goes offline

# Customer Information Collection Tool

The Customer Information Collection Tool is a utility feature used by Telestream Support and Development teams to collect information on the client's DIVA system to analyze and diagnose issues uncovered in the field. This utility is included in the DIVA delivery, but is only intended to be used by Telestream personnel.

The tool receives all customer information required for support investigations including log files, dump files, and client environment information. It receives information from all client sites in a uniform manner, and retains detailed client issue information with the originator's contact information. The tool also notifies the Telestream Development Team as soon as information is transferred to the development facility, where it is stored permanently for future issue resolution as necessary.

The CollectSysInfo.bat file enables collecting the required information to send to the Telestream Support and Telestream Development teams for issue resolution. Execute the batch file using the following command and parameters:

```
%DIVA_HOME%\Program\Utilities\bin\CollectSysInfo.bat [parameter
value]
```

Example:

```
%DIVA_HOME%\Program\Utilities\bin\CollectSysInfo.bat -EXMODULES
VACP, AMCommunicator -AFTERDATE 09/25/2016 -MACHINES
172.16.3.45,172.16.3.46 -DBTYPE conf -CUST -CUSTOMER1
```

The main CollectSysInfo.bat command parameters are as follows:

- `-EXMODULES [MODULE_NAMES]`

  Excludes the specified module from collection logs and configuration files. Using -`EXMODULES ALL` excludes all of the modules and only collect the Core Database dump. The default is collecting all modules.

- `-AFTERDATE [MM/DD/YYYY]`

  Collects logs only on or after the specified date. The default is collecting all available logs.

- `-MACHINES [IP:host_name,IP:host_name,and so on]`

  Collects the logs from any additional computers identified. Multiple host names are identified in a comma separated list. The default is to only collect logs for the current system where the script is running.

- `-DBTYPE [FULL|CONF]`

  Collects a full Core Database dump, or just a configuration dump. The default is collecting a full database dump.

- `-CUST [CUSTOMER_NAME]`

  The name of the customer where the logs are collected. The customer name is truncated if it is longer than 13 characters. There is no default value for this optional parameter. If it is not supplied as an argument the script prompts you to enter the Customer Name during execution.

There are also several internal parameters for the script. Each of the internal parameters has a default value that can be overridden by specifying a custom value using the script options.

**Example:**

```
%DIVA_HOME%\Program\Utilities\bin\CollectSysInfo.bat -EXMODULES
VACP, AMCommunicator -AFTERDATE 09/25/2016 -MACHINES
172.16.3.45,172.16.3.46 -DBTYPE conf -CUST CUSTOMER1 -DIVALOC
C:\INSTALL\DIVA
```

The additional script parameters are as follows:

- `-DIVALOC`

  The DIVA installation path for all computers from where the script is collecting logs. The default value is %DIVA_HOME%.

- `-REMOTEDIVA`

  The DIVA installation location if additional computers are specified using the `-MACHINES` parameter. The path set in this parameter must be shared within the network. The default value is \\RemoteSystem\C$\DIVA.

- `-DUMPPATH`

  The location where the script generates and outputs the .7z zip file. The default value is H:\.

- `-POSTGRESLOGIN`

  The Core Postgres Database user name and its connection details.

- `-CYGWIN`

  The Cygwin installation path. Default: `C:\cygwin\bin`.

- `-SEVENZIP`

  The 7z zip tool installation path. Default: `C:\Program Files\7-Zip\7z.exe`.

- `-TEMPDIR`

  The temporary directory where the script copies the logs and configuration files. This folder is created automatically at the beginning of the script execution and subsequently deleted after the script completes execution. The script fails execution if the path set in this parameter already exists. Default: `H:\supportinfo`.

# VACP Converter (Optional)

VACP (Video Archive Command Protocol) is a protocol developed by Harris Automation for interfacing to an archive system. DIVA has its own API for communicating with DIVA, which is not compatible with VACP.

To provide interoperability without the need to redevelop the archive interface at the automation level, this module is provided to act as an interface to convert VACP commands from the attached automation system to API commands.

The service requires a successful connection to DIVA to start. Therefore, it must be started manually either through the Windows Services component or from the command line after DIVA is running.

Use the following commands to start the VACP Service from the command prompt:

1. Open a Windows command prompt.
2. To start the VACP Service, enter *net start "VACP Converter"* at the command prompt. The quotation marks are required for services with spaces in their service name.
3. To stop the VACP Service, enter *net stop "VACP Converter"* at the command prompt. The quotation marks are required for services with spaces in their service name.

The VACPService.exe file enables you to run the VACP Converter as a service. Execute the file using the following command and parameters:

```
%DIVA_HOME%\Program\VACP\VACPService.exe command [options]
```

Appending the `-conf` (or `-f`) option after one of the following commands specifies a specific configuration file to load settings from. The VACPService.exe command parameters are as follows:

- `install` (`-i`)

  Installs the VACP module as a system service.
- `uninstall` (`-u`)

  Removes the VACP module service.
- `debug` (`-d`)

  Starts the VACP module in console mode.
- `version` (`-v`)

  Displays the version information and then exits.
- `help` (`-h`)

  Displays help information and then exits.

# Avid Connectivity (Optional)

Avid Connectivity with DIVA transfers archival data to and from DIVA in specific video formats and enables archiving and retrieval of single clips, or a sequence of clips. Avid Connectivity is no longer packaged with DIVA and is a separate installation process. Additional installation is required for certain plugins for both AMC and TMC.

All operations for the AM Communicator are performed from Avid Interplay, not DIVA. All TM Communicator archive operations are performed from Avid, while all restore and delete operations are performed from DIVA.

DIVA includes support for the Avid Web Services API for Archive, Restore, and Partial File Restore of clips and sequences directly from Interplay. Also included is AMC support for Interplay 3.8 and TMC support for Interplay 3.7 and 3.8.

Certain API operations used in Avid Connectivity (such as *GetByFilename* and *DeleteByFilename*) are not supported for complex objects.

See the Avid Connectivity and Tool book, or contact Telestream (see *Telestream Contact Information*) for more detailed information.

# SPM (Storage Policy Manager—Optional)

SPM (Storage Policy Manager) provides automatic migration and life-cycling of material within the archive based on the rules and policies defined in the SPM configuration. The SPM component is also used to trigger deletion of material from SPM managed arrays (based on disk space watermarks).

See the Storage Policy Manager User Guide (https://www.telestream.net/telestream-support/diva/support.htm) for detailed information.

# Miscellaneous Utilities

DIVA includes various miscellaneous utilities, some of which are associated with the modules previously listed in this chapter. The included utilities are as follows:

**%DIVA_HOME%\Program\Utilities\bin\DIVAConfigurationPrinter.bat**

Prints the current DIVA configuration. There are no command-line parameters.

**%DIVA_HOME%\Program\Utilities\bin\DivaScript.exe**

This utility enables using command line orders to execute jobs and operations.

**%DIVA_HOME%\Program\Utilities\bin\GetVersion.exe [app_path]**

Returns the version number for a specific application. The app_path is the valid path to the application being checked.

**%DIVA_HOME%\Program\Utilities\bin\rdtu.bat**

The RDTU—Recover Damaged Tape Utility—recovers object instances that reside on a damaged tape. The utility can recover instances that have valid copies on other available media (that is, internal tape or a connected disk array) within a local or remote DIVA system. There are no command line parameters. The settings and configurations are defined in the rdtu-conf.xml configuration file.

# Starting and Stopping DIVA

Starting and stopping DIVA involves specific, ordered processes.

## Topics

- Starting Content Conductor
- Stopping Content Conductor
- Content Conductor Failover Procedures

# Starting DIVA

To start the DIVA system, start the hardware first, then start the software in the sequence as described in the following sections.

## Starting DIVA Hardware

Perform the following steps in sequence to start all of the DIVA hardware components. Wait for initialization of each hardware component to complete before moving to the following step.

1. Confirm that all required devices are installed. If they are not installed, they must be installed before proceeding any further.

   **a.** Managed Storage and Drives

   **b.** SAN RAID Arrays

   **c.** Fiber Channel Switches

   **d.** Networking Devices

   **e.** Terminal Concentrator

   **f.** Graphical Front End Hosts (web app)

   **g.** Library DIVA Host

   **h.** External Direct Attached Devices

   **i.** DIVA Hosts

   **j.** Actor Hosts

2. Power on the Managed Storage and Drives.
3. Power on the SAN RAID Arrays.
4. Power on the Fiber Channel Switches (if installed).
5. Power on the Networking Devices.
6. Power on the Terminal Concentrator (if installed).
7. Power on the Graphical Front End Hosts (web app).
8. Power on the Library DIVA Host (if installed).
9. Power on External Direct Attached Devices.
10. Power on DIVA Hosts.

    In installations where two DIVA Hosts are installed, it may be required to always start the Main DIVA first, and then the Alternate (or Backup) DIVA at a later time. Consult with your Telestream Installer to determine if this is applicable to the installation.

11. Power on the Actor Hosts.

Hardware start is complete if everything powered on successfully.

## Starting DIVA Software

The following steps describe the required order that the software components of a DIVA system must be launched. Some software components may be set to launch automatically when the host is started (for example, the Actor Service).

**Note:** Some DIVA Windows Services may be disabled, or set to be launched in manual mode, due to the configuration and settings done by Professional Services during the installation.

The management of each DIVA software component, whether manually or automatically initiated, is covered in *Components*. Perform the following steps in sequence to start all of the DIVA software components:

1. Confirm that all required components are installed. If they are not installed, they must be installed before proceeding any further.
   - **a.** Library Robot Manager
   - **b.** Library
   - **c.** DIVA
   - **d.** Backup Service
   - **e.** Complex Objects (if in use)
   - **f.** DIVA Connect
   - **g.** VACP Converter
   - **h.** SPM (Storage Policy Manager)
   - **i.** WFM (Watch Folder Monitor)
2. Launch the Library Control software.
   - **a.** ACSLS
   - **b.** PCS
   - **c.** SDLC
3. Launch the Robot Cores.
4. Launch the Actors.
5. Launch DIVA.
6. Launch DIVA Connect.
7. Launch the VACP Converter.
8. Launch SPM.
9. Launch WFM.

Software start is complete if everything initialized successfully.

# Stopping DIVA

DIVA is stopped in the reverse order from starting the system. Shut down the software first and then the hardware. The following sections describe the required procedure to fully shut down DIVA.

## Shutting Down the Software

To ensure that jobs currently still in progress are not prematurely terminated by shutting down the DIVA system, it is recommended the DIVA be stopped first, because any jobs currently active are completed before DIVA completes shut down.

See *Stopping DIVA* for DIVA shut down procedures. When the DIVA is shut down all operations are ceased. It is not necessary to stop other software components before shutting down the host computer where they are installed.

## Shutting Down the Hardware

Use the following procedure (in sequence) to shut down all DIVA-related equipment and devices:

1. Shut down the DIVA Host.
2. Shut down the Actor Hosts.
3. Power off all External Direct Attached Devices.
4. Power off Graphical Front End Hosts.
5. Power off Terminal Concentrator (if installed).
6. Shut down the Library Manager Host (if installed)
7. Power off Network Devices.
8. Power off Fiber Channel Switches (if installed).
9. Power off SAN RAID Arrays (if installed).
10. Power off Library and Drives.

Hardware shut down is complete if everything powered off successfully.

# DIVA Failover Procedures

**Caution: These procedures are critical and sensitive. They should only be performed under the control of a Telestream Support Technician.**

The following steps are required to fail over a DIVA to the Backup when the database is still accessible on the original DIVA:

## Scenario 1—Failover with Multiple BKS Installations (Recommended)

Use the following procedure to fail over using multiple BKS installations. This is the recommended scenario. See the DIVA Database and BKS Installation, Configuration and Operations book for configuration details.

1. Stop all services on the Primary site (if available).
2. Start the BKS and DBAgent and related databases on the Failover site.

   If backups are enabled on the fail over site, you should disable them by setting the Enabled flag in the configuration file. This ensures that there are no backup operations competing with your failover command.

```
"DatabaseBackup": {
  "Enabled": false, <=== UPDATE
  "FullBackupInterval": {
    "executionPeriod": 0,
    "timeOfDay": "12:00:00",
    "instancesInPeriod": null
  },
  "IncrementalPeriod": 15,
  "FullBackupFileRetention": 10,
  "FullBackupArchiveRetention": 30,
  "ArchiveMediaGroup": "",
  "PermanentRetentionPeriod": 180,
  "ArchiveSourceName": "",
  "BackupExecutionTimeout": 120,
  "RestoreExecutionTimeout": 120,
  "StatusPollingPeriod": 3,
  "StatusReportingInterval": 1440
}
```

3. Initiate the failover command (This must be done for each database):

   **From the API: https:\\localhost:1877:**

   **a.** Select `PUT /Backup/failover/{name}` and click Try It Out.'

   **b.** Enter the database profile name for the source.

   **c.** Enter the failover profile name as the target.

   **d.** Enter a date if you are recovering from a specific time, or leave the field empty if you would like the latest backup file used.

   **e.** Click Execute.

**Note:** The API does not wait for status before returning. To see the status of the request you can use the `GET /Backup/status/{name}` endpoint with the name of the target profile.

**From the Initiator.exe:**

**a.** Select the Failover option.

**b.** Select the option for your failover. It looks similar to `<profile name> ? <profile name failover>`.

**c.** Select a time to failover from.

**d.** Wait for the operation to complete.

After the databases have been failed over, start other services and verify they are running as expected.

Enable the database backups in the configuration file that you disabled in an earlier step.

**Note:** To fail back, the same steps are run replacing the source and target database profiles. No files need to be transferred, all of this is done by the BKS.

## Scenario 2—Failover from a Single BKS Instance

Use this procedure to failover using a single BKS installation. See the DIVA Database and BKS Installation, Configuration and Operations book for configuration details.

1. Verify the DBAgent and the databases needed are up on the Failover server and that all services that use the database are offline.
2. Disable backups.

   If backups are enabled on the failover site, you should disable them by setting the Enabled flag in the configuration file. This ensures that there are no backup operations competing with your failover command.
3. Remove the primary databases from the managed primary location and add the failover databases to the managed databases for the failover location.

```
"LocationSettings": {
  "Locations": [
    {
      "Name": "Primary",
      "Primary": true,
      "Enabled": true,
      "Location": "H:\\divaback",
      "AgentUrl": "https://localhost:1878/",
      "Type": "Local",
      "ManagedDatabases": [], <=== UPDATE
      "BackupReplication": [
        "MetadataDatabaseFailover",
        "OracleDatabaseFailover"
      ],
      "SourceName": "",
      "User": "",
      "Password": ""
    },
    {
      "Name": "Secondary",
      "Primary": false,
      "Enabled": true,
      "Location": "\\<path to divaback on failover>",
      "AgentUrl": "https://<failover server ip>:1878/",
      "Type": "Local",
      "ManagedDatabases": [
        "MetadataDatabaseFailover", <=== UPDATE
        "OracleDatabaseFailover"    <=== UPDATE
      ],
      "BackupReplication": [
        "MetadataDatabase",
        "OracleDatabase"
      ],
      "SourceName": "",
      "User": "",
      "Password": ""
    }
  ]
}
```

**4.** Initiate the failover from the API or Initiator.exe:

**From the API: https:\\localhost:1877:**

**a.** Select `PUT /Backup/failover/{name}` and click Try It Out.'

**b.** Enter the database profile name for the source.

**c.** Enter the failover profile name as the target.

**d.** Enter a date if you are recovering from a specific time, or leave the field empty if you would like the latest backup file used.

**e.** Click `Execute`.

**Note:** The API does not wait for status before returning. To see the status of the request you can use the `GET /Backup/status/{name}` endpoint with the name of the target profile.

**From the Initiator.exe:**

**a.** Select the Failover option.

**b.** Select the option for your failover. It looks similar to `<profile name> ? <profile name failover>`.

**c.** Select a time to failover from.

**d.** Wait for the operation to complete.

**5.** Enable the database backups in the configuration file that you disabled earlier.

**Note:** To fail back, perform the above setups replacing the source and target databases and change which databases are being managed.

## Cluster Failovers

**Note:** Clusters are not supported in this DIVA release.

Use the following procedures if a cluster fails to initiate:

1. Check that the backups are synced on the Active Node and Backup DIVA Core.
2. Stop all DIVA Services from the Microsoft Cluster on the Active Node.
3. On the Active Node, run `SELECT COUNT(*) AO_OBJECT_NAME from DP_ARCHIVED_OBJECTS;`.
4. Create an export of the current database from the Active Node.
5. Stop the DB Services on the Active Node from the Microsoft Cluster Core.
6. Start the DB Services on the Backup DIVA Core server.
7. Recover the database from the backups. Contact Telestream (see *Telestream Contact Information*) if you require assistance recovering the database.
8. Start the DIVA Services on the Backup DIVA Core and run some tests to confirm functionality.

When all testing is successful, stop all Backup DIVA Core services, and restart all services from the Microsoft Cluster Core. Verify all operations are functioning normally.

# Using the DIVA Web App

The web app is a browser-based, platform independent, software utility that connects to the Manager through the REST API to monitor, control, and supervise operations in DIVA. Multiple instances can be operated simultaneously from any computer that has TCP/IP connectivity to both DIVA and the DIVA database.

The web app is not intended for the intensive archive operations of a DIVA system. Typically, archive operations are initiated to DIVA from a Broadcast Automation or Media Asset Management system. The web app is intended to supplement these operations rather than replace them.

The web app provides the following features:

- Monitoring of the jobs that have been submitted either through the Client API or from the web app.
- Monitoring the status of Actors, Drives, and Disks connected to DIVA.
- Initiate and submit all Client API available commands, such as Archive, Restore, Partial File Restore, and so on, to DIVA for execution.
- Management of tapes for each library controlled by DIVA (such as internalizing, externalizing, and tape repacking).
- Interrogation and data mining of the Core database.

See the *Step* section or contact Telestream (see *Telestream Contact Information*) for more information.

## Topics

# Launching the Web App and Connecting to DIVA

Use the following procedure to launch the web app and connect to DIVA:

1. Start the web app by opening a web browser and navigating to https://localhost:8765/DivaWebUI.

**Note:** Substitute "localhost" in the URL with the IP address of the Manager system if not directly on the main DIVA system.

2. The login window displays after the interface loads.
3. The dashboard is presented after logging in.

If the web app cannot establish a connection to DIVA, it displays a 404 error because the app is served by the web server in the Manager.

# User Permissions

After the connection to DIVA is established, the web app only allows access according to the profile permissions of the logged in user account. In order to obtain greater (or lesser) permissions, you must log out and then log back in using the appropriate profile. This table identifies the permission given to each specific profile:

| | Operator | AdvOperator | Admin | SysAdmin |
|---|---|---|---|---|
| Restore Object | | X | X | X |
| Archive Object | X | X | X | X |
| Copy Object | X | X | X | X |
| Stage Object | | | X | X |
| Delete | | | X | X |
| Transfer Files | | X | X | X |
| Insert Tape | | X | X | X |
| Eject Tape | | X | X | X |
| Repack Tape | | X | X | X |
| Import Tape | | X | X | X |
| Export Tape | | X | X | X |
| Clone Tape | | | X | X |
| Verify Tape | | X | X | X |
| Assign Storage Plan | | | X | X |

| | Operator | AdvOperator | Admin | SysAdmin |
|---|---|---|---|---|
| Protect Tape | | | X | X |
| Edit Tapes | | X | X | X |
| Associate Sets | | X | X | X |
| Auto-Repack Tape | | | X | X |
| User Management | | | | X |
| General Settings | | | X | X |
| Configuration[1] | X | X | X | X |
| Drive Update | | X | X | X |
| Disk Update | | X | X | X |
| Notify Manager | | X | X | X |
| SPM Actions | | | X | X |
| Migration Tool | | | X | X |
| Cloud Bucket Scanning | | | X | X |
| Media Priority Update | | X | X | X |

[1] Operators and AdvOperators have limited permissions in the Configuration area.

# Web App Preferences

Web app preferences are controlled by the Configuration > Local Settings page. Date and Time display preferences, Home Page selection, and Banner Shortcuts selection.

All settings are stored locally in the browser settings and do not follow the logged in user on other systems. The local settings can be reset to defaults on the Configuration > Local Settings > Reset Local Config page. The reset can apply to all pages using the slider button at the top of the page (Apply To All), or only to specifically selected pages using the slider buttons next to each page.

All pages with tables remember the selection of which columns are displayed in a table, and in what order, column width, and page size.

# Dashboard and Quick Launch Buttons

The web app dashboard displays information at a glance as soon as the application is started and you are logged in; it is designed to give a quick overview of the system and it's usage.

The dashboard home page displays four panels;

- System Info
- Utilization Stats
- Alerts
- Storage Resource

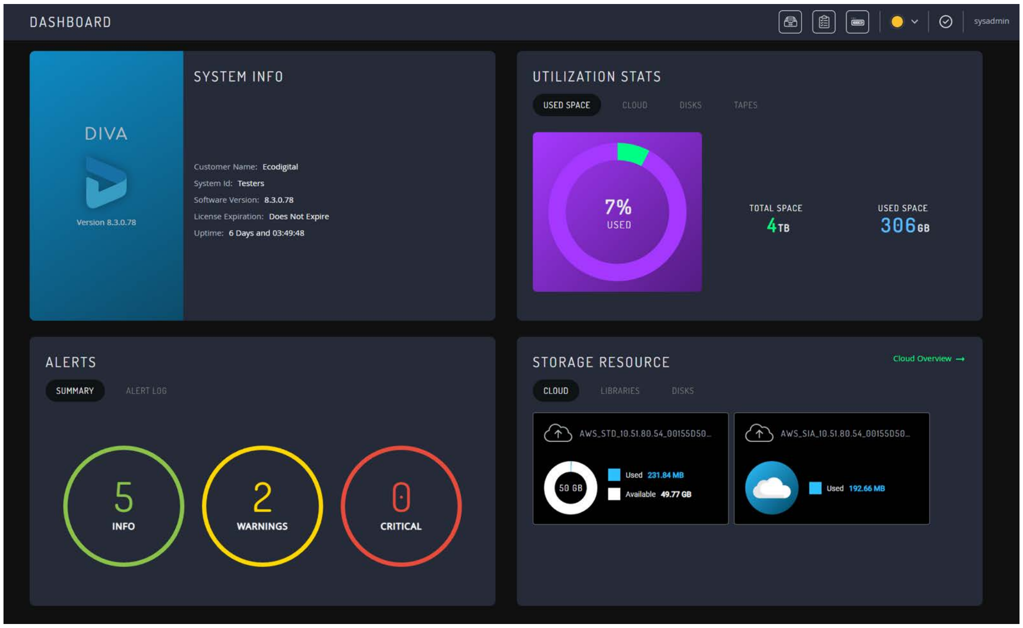

Version 8.3.0.78

Customer Name: Ecodigital
System Id: Testers
Software Version: 8.3.0.78
License Expiration: Does Not Expire
Uptime: 6 Days and 03:49:48

7% USED
TOTAL SPACE 4 TB
USED SPACE 306 GB

5 INFO
2 WARNINGS
0 CRITICAL

AWS_STD_10.51.80.54_00155D50...
50 GB
Used 231.84 MB
Available 49.77 GB

AWS_SIA_10.51.80.54_00155D50...
Used 192.66 MB

The System Info panel displays some properties from the license. For example, Customer Name, System Id, and the License Expiry Date. It also displays the software version of DIVA that is running, and how long the Manager has been running (uptime).

The Alerts panel displays a summary of how many of the metrics it is tracking are Info, how many are Warnings, and how many are Critical. The details of these metrics are not on this panel. The metrics allow a quick view as to whether all Drives, Libraries, Disks, Cloud Disks, and Actors are available.

Click the Alert Log link at the top of the Alerts area to display the status of DIVA components and shown here:

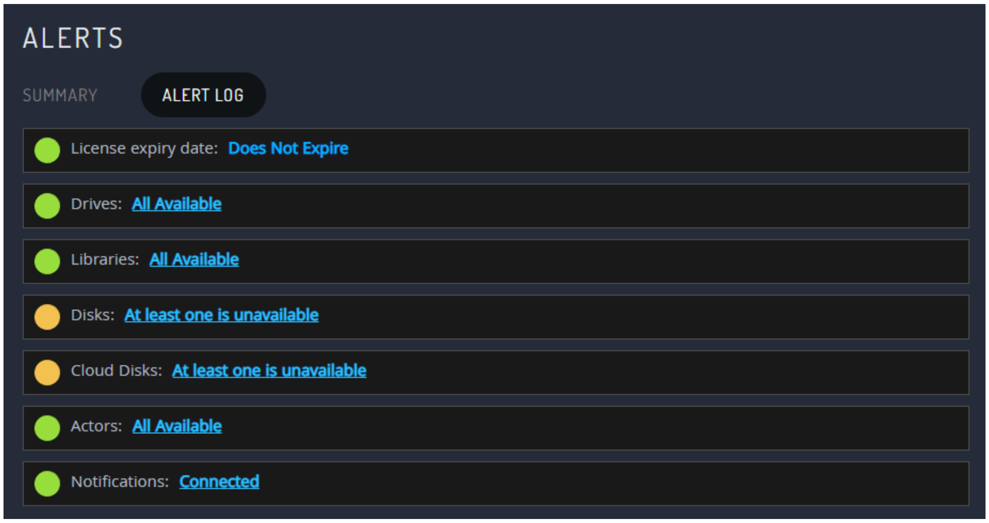

License expiry date: Does Not Expire

Drives: All Available

Libraries: All Available

Disks: At least one is unavailable

Cloud Disks: At least one is unavailable

Actors: All Available

Notifications: Connected

The Utilization Stats panel displays the total space available to the system for storage, and how much of that space is currently consumed. There is a combined view of Cloud, Disk, and Tape, and also these metrics broken out in to their individual categories.

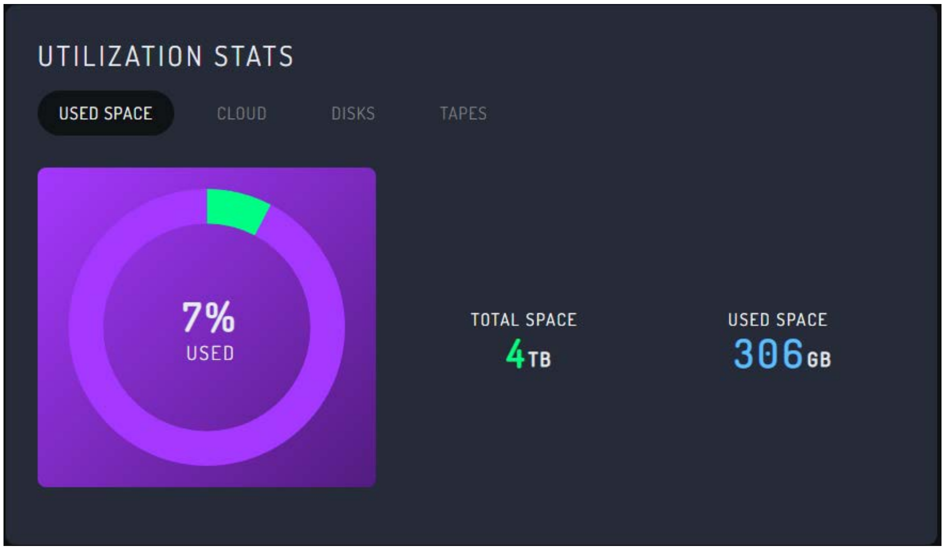

TOTAL SPACE 4TB

USED SPACE 306GB

Because Cloud storage is technically unlimited, this is omitted from the combined view unless a quota limit on the Cloud Bucket has been set in the configuration. The quota limit is only for display purposes on the dashboard. If it is not set, then the Cloud Disk displays just the amount of used space. The Cloud Quota is set on the Configuration > Cloud Storage > Cloud Buckets > Add Cloud Bucket (or the Edit Cloud Buckets) page.

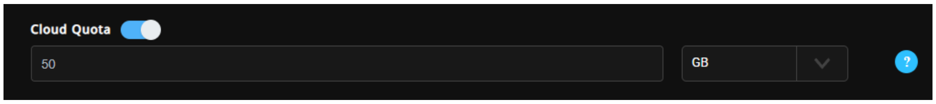

Cloud Quota: 50 GB

The Storage Resource panel displays similar information to the Utilization Stats panel. However, rather than totaling the information for all resources in a category, it breaks them down in to individual resources. Click on the Cloud, Libraries, or Disks links at the top of the Storage Resource panel to display information for that specific resource.

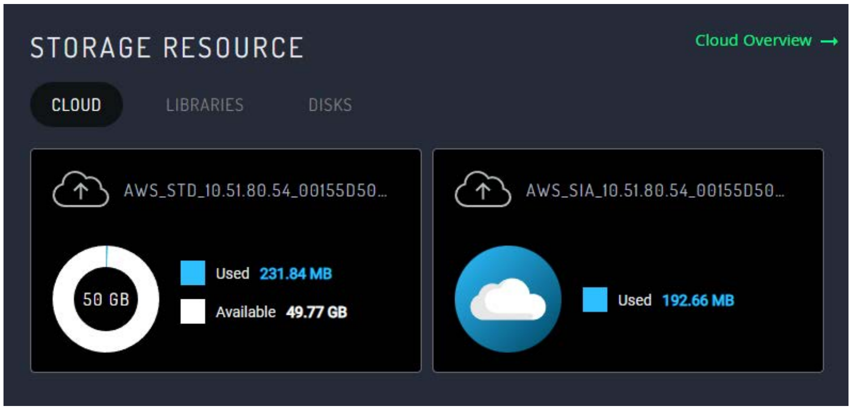

STORAGE RESOURCE

AWS_STD_10.51.80.54_00155D50...
50 GB
Used 231.84 MB
Available 49.77 GB

AWS_SIA_10.51.80.54_00155D50...
Used 192.66 MB

Users can change the default home page from System Dashboard to Jobs in the Configuration > Local Settings > Customization page using the pull-down list.

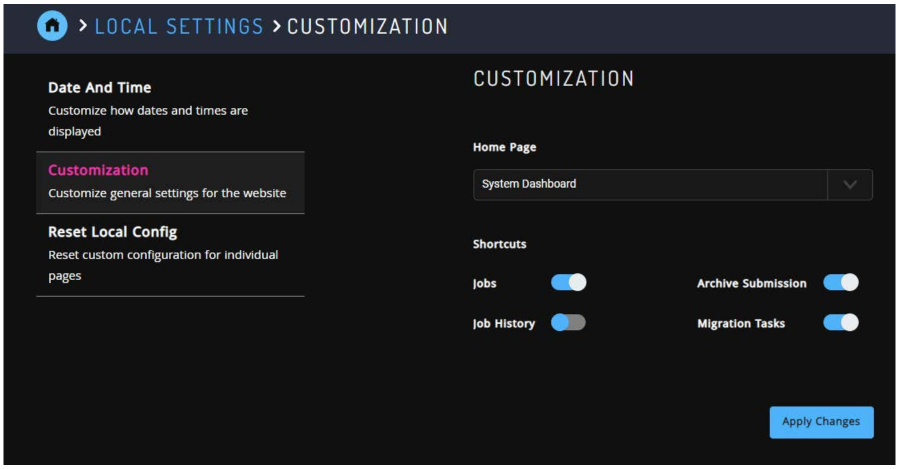

> LOCAL SETTINGS > CUSTOMIZATION

**Date And Time**
Customize how dates and times are displayed

**Customization**
Customize general settings for the website

**Reset Local Config**
Reset custom configuration for individual pages

CUSTOMIZATION

Home Page: System Dashboard

Shortcuts
Jobs: on
Archive Submission: on
Job History: off
Migration Tasks: on

Some page shortcuts for regularly visited pages are located on the top banner for easier navigation. Users can turn on or off the shortcuts from the Configuration > Local Settings > Customization page using the slider buttons to enable or disable the shortcut. There are shortcuts for Jobs, Job History, Archive Submission, and Migration Tasks.

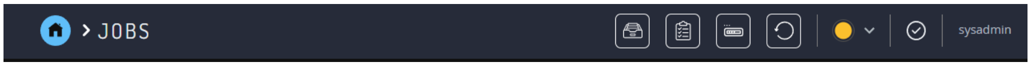

## Archive Submissions

The Archive Submission button, which opened a form, has been removed from the Catalog Browsing page, and a new dedicated page has been added for this action.

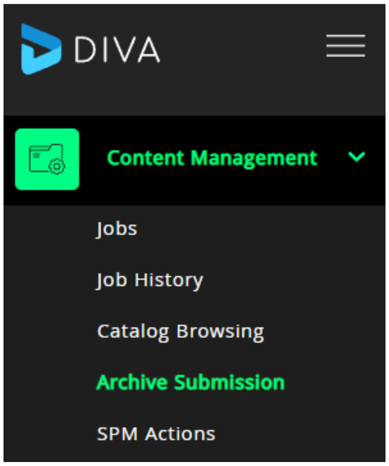

The original form is replaced with a dedicated page so it is less cramped and fits the required controls. It also allows the page to be in the main Content Management menu on the left, making it easier and quicker to find:

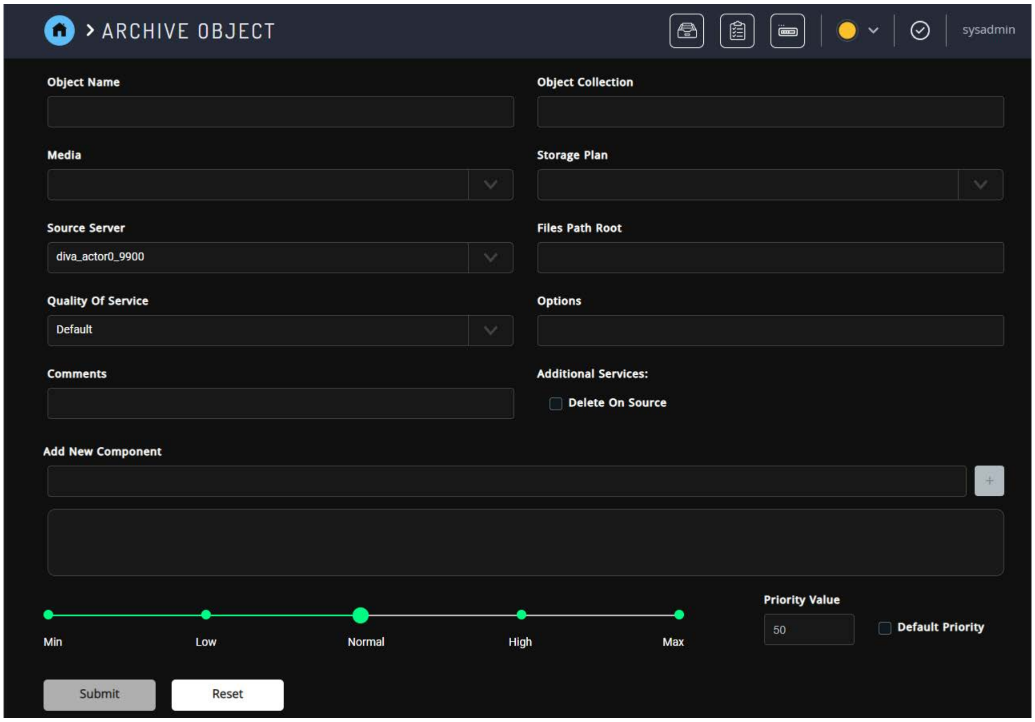

Source Server: diva_actor0_9900

Quality Of Service: Default

Priority Value: 50

# Toolbars and Navigation

Various functions are performed from the menu on the left of the screen.

The menu contains the following items:

- Content Management
  - Jobs
  - Job History
  - Catalog Browsing
  - Archive Submission
  - SPM Actions
- Resources Management
  - Unmanaged Storage
  - Actors
  - Media
  - Disks
  - Cloud Disks
  - Libraries
  - Drives
  - Tapes
- Troubleshooting
  - Logged Events
  - Drive Alert Logs
  - Library Alert Logs
- Migrations
  - Start New Migration Task
  - Migration Tasks
- Configuration
  - User Management
  - System Settings
  - Library Storage
  - Cloud Storage
  - Disk Storage
  - General Settings
  - Local Settings
  - License

The web app release information is displayed in the System Info panel on the dashboard.

To display the web app connection information, click the Alert Log link in the Alerts panel of the dashboard.

The following sections describe different web app areas.

**Note:** Clicking any of the panel views displays the Properties Page for that device; otherwise tables have a link of the Object Name.
All table views include a download button on the top right of the page. Click the download button to download the current table in CSV format to the local machine.

## Dashboard

Click the Start orb or the Home icon (displays when viewing pages other than the dashboard) on the top left to display the web app main dashboard, which displays general system information and statistics. The dashboard view is described in the previous section.

## Jobs and Job History

The Jobs view primarily displays the current jobs submitted to DIVA that are currently in, or pending, execution. Jobs that have completed, canceled, or encountered warnings during execution are displayed on the Jobs History page. This feature only applies when the web app is connected to DIVA. Completed or canceled jobs before the connection are not displayed. While connected, the number of pending, executing, completed, and canceled jobs displayed is a maximum of 300 of the most recent jobs. Clicking a job produces an informational screen pertaining to the selected job.

### Step

The Step column indicates the current operation of the job being performed by DIVA. A short description of each step is as follows:

### Mounting

A tape is being inserted into a drive. The mounting step is completed after the tape is fully threaded, positioned at the tape header, and the DIVA label on tape verified with that of the barcode label.

If the label does not match, the job is aborted, and the tape is set to Not Writable. This situation could occur if there is a mismatch in the Actor-Drives configuration, or the tape has already been used in another archive system (and therefore has a tape label not usable by DIVA). The latter example is a protection feature in shared library environments, where the tape has not been set to Set ID 99 (that is, not in use by DIVA).

Tapes from other archive systems must first have their tape label cleared before using them in DIVA. Contact Telestream (see *Telestream Contact Information*) for assistance to use such tapes.

## Dismounting

This step involves ejecting a tape from a drive. An Actor first issues an Eject command to the drive (in which it rewinds, unthreads and ejects the tape), and the Robot Manager issues a dismount command to the library to return it to a tape bin. If the drive cannot complete the job, the job is terminated and the drive set to Out of Order.

## Positioning

When reading from tape, the tape is positioned to the selected object. When writing to tape, the tape is positioned to the End of Data (that is, to the position where the last object was written). If this process takes too long, DIVA times out the operation and attempt to dismount the tape. If this also fails, the drive is set to Out of Order.

## Reading

The object displayed in the Virtual Object column is being read from tape. If this step takes too long (for example, the drive is in a hung state), DIVA times out the step and attempt to use another drive (or instance if available).

## Writing

The object displayed in the Object Name column is being written to tape. If this step takes too long, (for example, the drive is in a hung state), DIVA attempts to write the object to another tape in another drive.

## Deleting

This step is rewriting the tape's label and moving the End of Data pointer therein before writing to the tape in an archive operation. This is seen on a tape which has previously been used by DIVA but has since had all objects deleted and consequently returned to the Unused Tapes Sets.

## Transferring

Data is either being transferred to a Source Server from an Actor cache or from a destination to the Actor cache.

## Waiting for Operator

This step holds the job in a suspended state and is waiting for human intervention, such as inserting tapes in library's Cartridge Access Port.

## Waiting for Resources

DIVA resources required for this job are currently in use by another job and the job is executed when they become available. Resource availability can also be influenced by the Priority of other jobs much lower in the queue.

DIVA incorporates an intelligent feature (which must be enabled in the DIVA configuration) whereby jobs with a lower priority are boosted over higher priority jobs if they involve a tape that is already mounted from a previous job. This feature can substantially reduce the amount of tape mounting and dismounting, and speed up the executions of all jobs overall.

## Clearing Completed Jobs

Completed, Aborted, Partially Aborted, or Canceled jobs can be cleared from the Current Jobs Queue by clicking the Clear or Clear All buttons in the view or the Current Jobs context menus.

## Canceling a Job

A running or pending job can be canceled by first selecting the specific job to be canceled, and either clicking Cancel from the command menu or from the Current Jobs context menu.

**Note:** The current operation (or step) on a job that is currently being executed may need to complete before the job is actually canceled by DIVA.

## Changing the Job Priority

If there are several pending jobs in the Current Jobs Queue, DIVA processes each job based on its Priority.

If a specific job needs to be executed before (or after) the jobs that precede it in the queue, the job's priority can be manually adjusted to be higher (or lower) than that of the preceding jobs. Raising (or lowering) a job's priority can also be achieved through the Client API using a third party archive initiator.

Increasing a pending job's priority does not stop, or pause, any jobs that are currently executing. This simply changes the order in which the pending job is processed by DIVA, except if a resource being used by a running job (such as a specific tape) becomes available after that running job has completed an operation. The order of job execution may also be influenced if the `DIVA_CORE_PRIORITY_TIER` setting in the DIVA configuration is enabled (that is, a job lower in the queue has this value added to its priority if it involves a tape already mounted).

By default, DIVA periodically increments the priority of all jobs in the queue. This prevents low priority jobs (such as Copy to Group) from continually being overridden by higher priority jobs and getting stuck indefinitely in the queue.

## Retrying a Job

A previously completed or failed job can be resubmitted to DIVA using the Retry command. This is useful for resubmitting similar jobs where few of the details change between them. For a job that was terminated (for example, because a parameter was incorrectly entered, or a server was briefly offline), the failed job can be retried without having to submit a completely new job.

## Actors

The Actors view provides an indication of the status of each Actor defined in the configuration and any current jobs. This view is displayed by navigating to Resources Management > Actors in the left menu.

Select one of the Actors to displays the Actor Properties panel. If the DIVA cannot establish a connection to an Actor, it is displayed as Offline. Click an Actor Name, or the three dots next to the Actor Name and select Properties to display that Actor's Properties page.

The top right of the Actors page includes a button to change the view between the table view and item view. Use the button to switch views back and forth. The last view you were viewing is retained and is the initial view when you log out and back in.

## Libraries

The Resources Management > Libraries view gives the information and an indication of the status of each of the libraries connected to DIVA. This screen displays the Serial Number, Name, Type, ACS, Status, First Utilization Date, Total Tapes, Total Data Stored, Total Capacity, and Free Capacity for the connected Libraries.

Click the three dots (on the right) in the library area to display the Properties screen. This view offers information concerning the specified library.

## Drives

The Resources Management > Drives view displays the status of each tape drive in the libraries connected to DIVA, what (if any) tape is mounted in each drive, and current operations being performed on the tape in the drive. Online or offline status for a drive is configured in the Resources Management > Drives menu item.

If DIVA encounters a problem with a particular drive, it sets the drive to Out of Order as a safety measure. When a drive is set to this state it is not used for DIVA operations.

The top right of the Drives page includes a button to change the view between the table view and item view. Use the button to switch views back and forth. The last view a user was using is retained and is the initial view when you log out and back in to the system.

**Note:** If a drive has been set to Out of Order, the cause of the error must be investigated before setting the drive back to Online.

# Disks

The Resources Management > Disks view displays the online status and capacities of disks configured in DIVA. The status of a disk can be set through the Disks page using the three dots pull-down list and selecting the appropriate status. If DIVA has automatically set a disk to Out of Order, the cause of the error must be investigated before setting the disk back to Online. When DIVA encounters an I/O error with the disk, it is set to Out of Order automatically by DIVA.

The column titled Consumed Size represents the space in kilobytes consumed by the content on disk.

## Viewing Storage Options

The disk storage options can be viewed on the Resources Management > Disks screen. When a Disk Name is clicked, the Disk Properties page is displayed. Click the Files on Disk link (under the disk properties) to display a panel showing the files contained on the disk.

## Cloud Disks

The Resources Management > Cloud Disks view displays the online status and capacities of the cloud disks configured in DIVA. The status of a cloud disk can be set through the Cloud Disks page by clicking the three dots next to the Disk Name and selecting the appropriate status. If DIVA has automatically set a disk to Out of Order, the cause of the error must be investigated before setting the disk back to Online. If DIVA encounters an I/O error with the disk, it is set to Out of Order automatically by DIVA.

The column titled Consumed Size represents the space consumed by the content on disk. This column is especially useful for cloud accounts with unlimited disk space, because it provides visibility into the amount of content stored in the cloud.

**Note:** A disk linked to an OCI storage account always reports a Consumed Size of 0.

The Cloud Storage Options column associated with the array is also displayed in this view. OPC Cloud disks have a storage class of Standard (immediately available for download from the cloud) or Archive (requires a maximum four hours to download from the cloud). OCI cloud disks only support STANDARD storage.

## Viewing Storage Options

The disk storage options can be viewed on the Resources Management > Cloud Disks > Properties page by clicking on the Disk Name. A list of Scan Events is displayed at the bottom of the screen. Use the Actions pull-down menu to configure and start Cloud Bucket Scanning.

# Tapes

This view is displayed by selecting Resources Management > Tapes on the left menu. The Tapes view provides flexible search criteria (the Filter button located at the top left of the screen) to execute DIVA database queries about the tapes managed by DIVA.

On the Tapes page, either click the three dots to display a context menu, or select one or more check boxes for the tapes to perform various tape operations. Selecting one or more check boxes displays action buttons at the bottom of the screen.

Most tape operations are self-explanatory. However the Automatic Repack operation requires some description and is covered in the following section.

For Sony Optical Drives, whether a media is Write-Once can be viewed by clicking on the tape (the Write-Once property is displayed in the Tape Properties window). The Blu-ray discs are shown as tapes and viewable in the Tapes view. The Write Once Media column displays this information as either Y or N indicating whether the tape is Write-Once. The view can also be filtered so only Write-Once media is displayed.

Click the More menu for a specific tape in the Tapes view to use a menu with the options available on the selected tape. Click the Tape Name to display the Tape Properties page, which is for informational purposes only. No data within the tape can be directly manipulated by an operator from this page.

## Tape Compression

Tape compression is supported at the Tape Group level.

When tape compression is enabled, any empty tape assigned to the Tape Group have compression enabled, and instances written to the tape are compressed. Tapes assigned to the Tape Group before compression was enabled remain uncompressed, and instances written to the uncompressed tape are uncompressed.

To view all tapes with compression enabled, navigate to the Resources Management > Tapes page, open Filters, and set the Compression Enabled filter to Yes, then click Apply Filters.

## Tape Drive Encryption

Tape drive encryption securely supports bulk tape migration between DIVA systems. Tape Group level encryption is enabled, disabled, or updated in the Tape Groups view of the Resources Management > Tapes page using the Filter button.

See the Installation and Configuration Guide (https://www.telestream.net/telestream-support/diva/support.htm) for detailed configuration, and export and import information.

## Manual Tape Cloning

Users either select Clone Tape from the three dots next to the desired tape barcode and select Clone Tape from the pull-down menu, or select the checkbox for the tape to display the action button at the bottom of the screen. After one or more archives to a Tape Group containing a Clone Tape Group, users can manually clone the tape containing the newly archived objects. For example, upon archiving of six objects to a

Tape Group called *GroupA* configured with a Clone Tape Group named *GroupB*, a user can clone the tape containing the newly archived objects using one of the aforementioned methods.

Users can view the objects on the Source Tape and Clone Tapes before actually cloning the tape. Comparing the objects on each tape, the Clone Tape may be missing six objects. In this example, the Source Tape is also not synchronized with its clone. This can be seen by viewing the synchronized flag of the Clone Tape in the Tape List View. To exclusively copy the missing content from the Source Tape to the Destination Tape, click the Source Tape containing the newly archived objects, then select Clone Tape, and submit the job.

After selecting the Clone Tape menu option (or action button) the following pop-up form is displayed. Click Submit to clone the tape.

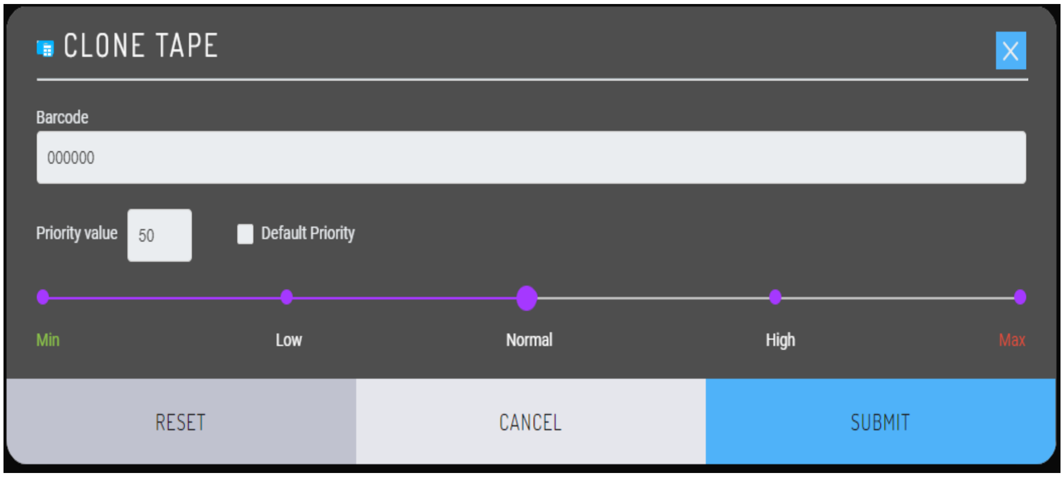

Barcode
000000

Priority value 50

□ Default Priority

During the first clone of a Source Tape, an empty tape contained in the set associated with the Clone Tape Group is selected. The selected tape is of the same type as the Source Tape, and equal to or greater in size. The newly selected Clone Tape is marked as protected.

Only clone jobs can write to a Clone Tape, and any attempts to write to the Clone Tape using any other job will terminate. Every subsequent clone of the same Source Tape writes to the same Clone Tape. The format of the Clone Tape is the same as that of the Source Tape.

Verify that the Clone Tape contains the same content as the Source Tape by viewing the Clone Tape after the clone job completes. The remaining space, number of elements, objects and positions all match the Source Tape.

Also, the Synchronized state of the Source Tape is updated to reflect that the Source Tape is now synchronized with its clone.

## Automated Cloning

In addition to manually cloning a tape, users can set up the periodic cloning of all Tape Groups configured for cloning by configuring two settings on the Configuration >

General Settings > Media page. The first setting regulates the frequency of the automated clones; the value identifies the number of hours between the submission of clone jobs. The second parameter determines the maximum number of simultaneous clones. In the following figure, a user has configured the periodic cloning of a maximum of 10 tapes every hour. Tapes are cloned in order of least recently cloned. A clone job is only attempted if Virtual Objects were written to the tape after the last clone.

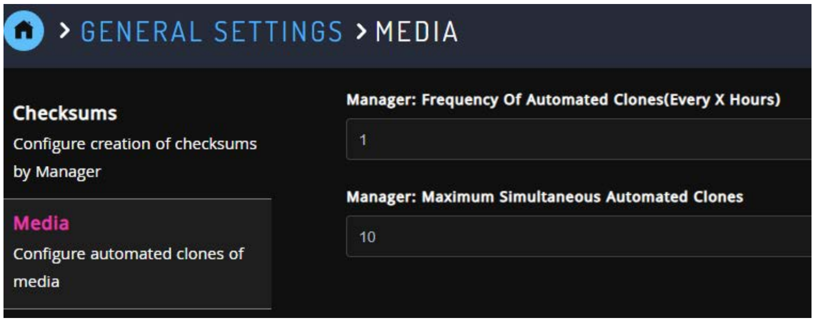

**Checksums**
Configure creation of checksums by Manager

**Media**
Configure automated clones of media

Manager: Frequency Of Automated Clones(Every X Hours)
1

Manager: Maximum Simultaneous Automated Clones
10

## Automatic and Manual Repack

When DIVA writes an object to a tape, the object can only be appended to where the last object was written on that tape. When any object is subsequently deleted from a tape, the space from that object cannot be reused. Eventually, as more and more objects are deleted, tape fragmentation occurs and potentially develops a considerable amount of unusable space in the tape library.

This unusable space can be reclaimed by repacking the tape. The repack process reads all material from the tape being repacked to a temporary cache and then writes it back to a new tape in the same Tape Group as the original (sequentially).

A manual tape repack can be performed by clicking the Repack Tape button, or triggered automatically when tape fragmentation and used capacity thresholds are exceeded. The following figure is displayed when performing a manual tape repack:

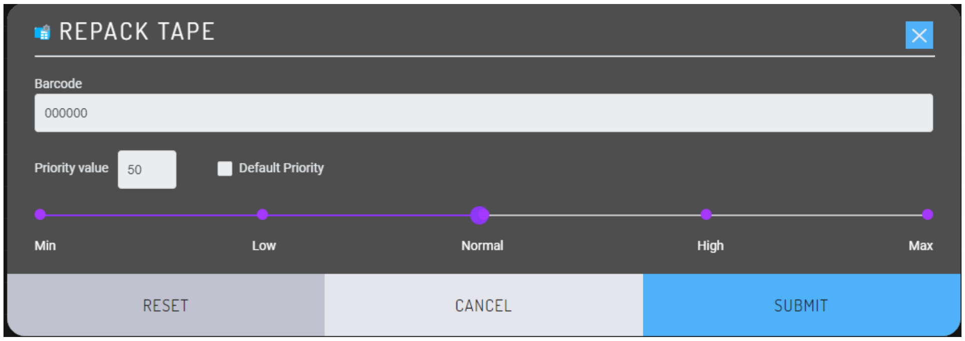

Barcode
000000

Priority value 50
□ Default Priority

Use the following procedure to enable or disable) Automatic Repack:

1. Navigate to Resources Management > Tapes.
2. Select the check box(es) next to the desired tapes.

3. Click the Automatic Repack button at the bottom of the tapes list to display the form.

---

**Note:** Automatic Repack ignores WORM Media. If WORM Media is repacked manually, the space is not recoverable.

---

4. To enable Automatic Repack select the check box, fill in the desired Tape Filling Threshold, Tape Fragmentation Threshold, Start Time, Duration, and select the groups to repack from:

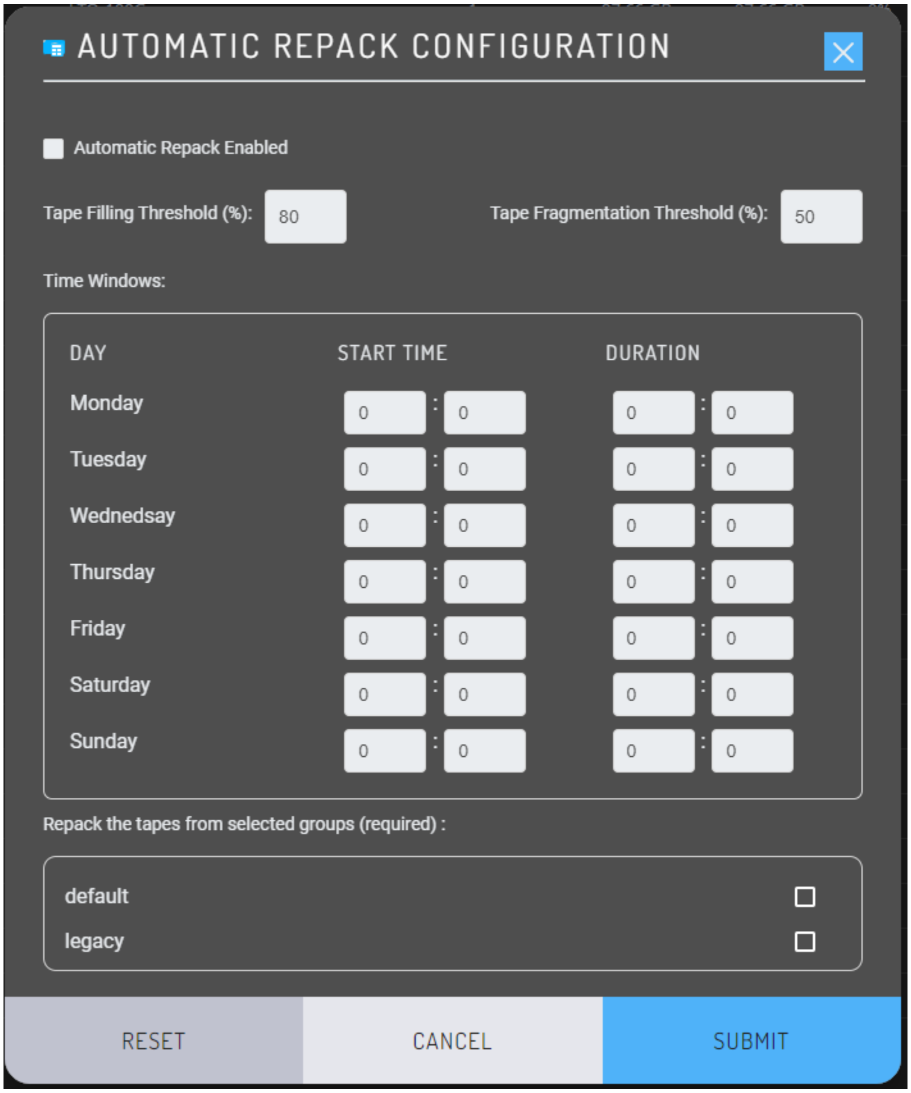

AUTOMATIC REPACK CONFIGURATION

□ Automatic Repack Enabled

Tape Filling Threshold (%): 80

Tape Fragmentation Threshold (%): 50

Time Windows:

| DAY | START TIME | DURATION |
|---|---|---|
| Monday | 0 : 0 | 0 : 0 |
| Tuesday | 0 : 0 | 0 : 0 |
| Wednesday | 0 : 0 | 0 : 0 |
| Thursday | 0 : 0 | 0 : 0 |
| Friday | 0 : 0 | 0 : 0 |
| Saturday | 0 : 0 | 0 : 0 |
| Sunday | 0 : 0 | 0 : 0 |

Repack the tapes from selected groups (required) :

default □
legacy □

Click the Submit button.

To disable Automatic Repack, select the tape check box, click the Automatic Repack button at the bottom of the tape list and deselect the check box; then click the Submit button.

# Unmanaged Storage

The web app Resources Management > Unmanaged Storage view provides information about the sources and destinations identified in the DIVA system. This view displays the Source or Destination Server Name, Product System, Type, Address, and First Utilization Date. Clicking one of the entries in the Name column displays the Properties page for that resource.

The information on the Properties page is not editable and is only for informational purposes. However, certain settings can be copied to the computer's clipboard by clicking on the double square boxes next to the property (for example, Root Path or Options)

# Objects

This view is displayed by clicking Content Management > Catalog Browsing in the left menu. Users can search objects using the Filter button. Users can also edit the content of the display by clicking the Edit Table icon on the top right of the screen.

Select a object and click the More menu to display the objects context menu. Select the check box next to an object to display these command as buttons at the bottom of the window. The form that displays when you click a button does not automatically specify an instance of the object in the Instance field of the job. You can manually select a specific instance using the pull-down list in this field before the command is sent. If the Instance field is left blank, the command deals with all instances of the object. For example, if Delete is selected, and no instance number is specified in the job, DIVA deletes all instances of the object.

Click the Object Name to display an Object Properties dialog box that includes the object's properties, instances, and components. A valid instance number must be specified in any command issued from this view or error message is displayed when submitting the command.

This includes removing the instance number entirely from the job. For example, all instances of an object cannot be deleted from this page by leaving the Instance field empty. This view doesn't permit the last instance of the object to be deleted and DIVA automatically terminates this job.

If a file (or part thereof) of an object instance is spanned across two or more tapes, and only one tape of the set is externalized, the instance is still considered externalized. However, an object (that is, all instances) is only considered externalized if all instances of that object are externalized.

Clicking the Content Management > Job History displays the Object's job history screen. Click the three dots next to an object, or the object ID, to display the job information for the selected object.

When restoring the same file to the same destination twice in parallel, the behavior on Windows and Linux is different. On Windows, the first restore (they cannot arrive exactly at the same time) locks the file so that the second one can terminate. On Linux, there is no such lock at the file system level. Both restores are executed at the same time, and both write to the same file. The content of the resultant file is not predictable.

## Jobs

This view is displayed by clicking on Content Management > Jobs, or Jobs History in the left menu. This view is limited to 300 lines by default. Completed, Canceled, and Aborted jobs are cleared if the web app is disconnected or relaunched. The Jobs view is provided for retrieval of previously completed jobs from the DIVA database. It is commonly used by Telestream Support to troubleshoot a previously reported issue.

After a particular job's ID is determined, it can be found in the Jobs History view. Users can search the displayed jobs using the Filter button at the top of the window.

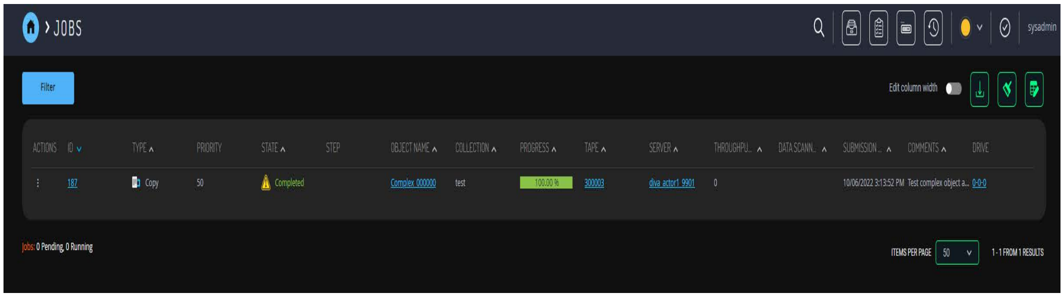

| ACTIONS | ID | TYPE | PRIORITY | STATE | STEP | OBJECT NAME | COLLECTION | PROGRESS | TAPE | SERVER | THROUGHPU... | DATA SCANN... | SUBMISSION ... | COMMENTS | DRIVE |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| ⋮ | 187 | Copy | 50 | Completed | | Complex 000000 | test | 100.00 % | 300003 | diva actor1 9901 | 0 | | 10/06/2022 3:13:52 PM | Test complex object a... | 0-0-0 |

Jobs: 0 Pending, 0 Running

ITEMS PER PAGE 50

1 - 1 FROM 1 RESULTS

A yellow triangle in the State column indicates the job completed but had a warning associated with it. A red triangle in the State column indicated the job failed with an error. Click the job ID to open the Job History > Properties screen to identify any errors or failure events. On the Properties page the cause of the failure is stated in the Abort Reason area on the right, and also in the Job Events area at the bottom of the screen.

The Job History page displays information (including Job Properties, Object Properties, Archive Properties and Events List) about the selected job. Up to fifty thousand logged jobs are stored in the DIVA database. After this number of jobs is reached, the oldest jobs are overwritten. For most facilities, this provides at least approximately six to twelve months of logged jobs.

## Media

This view is displayed by clicking on Resources Management > Media in the left menu This view displays information for each of the Tape Groups and Disk Arrays identified in the DIVA system. The search can be filtered using the Filter button at the top of the screen.

The Name list can be a full or partial media name including wildcards. Use an asterisk to display all media names.

Use the Type list to select viewing all media types, only Tape Groups, or only Arrays.

Click Refresh after making the filtering selections to update the display.

Clicking on the Tape Group or Disk Array displays the Media Properties page with details about that Tape Group or Array. The screen is for informational purposes only and is not editable.

### Source Media Priority

The Source Media Priority determines which source instance is preferred (according to the media where the instance resides) during the instance selection process of a Restore, Partial File Restore, and Copy To Group job. Instances on media with a higher priority are preferred. Cloud instances are only copied or restored if all local instances are offline, or no local instances exist. This is an absolute condition independent of the Source Media Priority.

## SPM Actions

This view is displayed by clicking on Content Management > SPM Actions in the left menu. This view is only applicable to installations having SPM (Storage Policy Manager) installed. It enables more detailed information to be extracted from the Core database related to the actions that have been initiated to DIVA from the SPM module.

Click the Filter button at the top of the screen to limit the type and number of results returned from the query. Run the query by clicking Apply Filters on the Filters popup dialog.

Clicking the three dots under Actions on a result returned by the SPM Actions query displays the SPM Actions Context menu. The menu has only two options as follows:

- Reschedule Action

  If the SPM-initiated job failed (for example, the medium or associated Actors for the slot were unavailable), this enables a retry of that SPM Action.

- Mark Action Completed

  Selecting this option from the context menu allows marking the SPM action as Complete if it is in another state; for example partially completed.

## Logged Events

This screen is displayed by navigating to Troubleshooting > Logged Events in the left menu. The Events view is typically used with the Job History view for troubleshooting purposes. The displayed results can be filtered using the Filter button at the top of the screen including Dates (Start and End dates and times), Severity (Information, Warnings, Errors, and Critical), Job ID, and Description.

When a particular job fails, the log of that job can be exported to a text file and sent to Telestream Support (when requested). This information can also be collected directly by the Telestream Support Engineer using the Customer Information Collection Tool.

When the query is run for the failed job's Job ID (usually retrieved from the Job History view), it shows the same events of that job's event log. This file can be saved as a text file by selecting Export.

DIVA stores a maximum of one million events in its database. When the number of logged events exceeds this value, DIVA begins overwriting the existing events beginning with the oldest entry.

## Drive Alert Logs

This screen is displayed by navigating to Troubleshooting > Drive Alert Logs in the left menu. This view lists errors reported by tape drives. This information is vendor-specific and may vary depending on the make and model. The search can be filtered using the Filter button at the top of the screen. For example, errors related to a particular tape can be filtered.

## Library Alert Logs

This screen is displayed by navigating to Troubleshooting > Library Alert Logs in the left menu. This view lists errors reported by direct-attached SCSI protocol Managed Storage. This information is vendor-specific and may vary depending on the library make and model. The search can be filtered using the Filter button at the top of the screen.

# Library Storage

DIVA Library Storage is in the Configuration >Library Storage menu to allow libraries to be configured in the web app by System Administrators and Administrators.

**Note:** Operators, Advanced Operators, and Users do not have access to the options described in this section.

# Robots

The default page displays the Robots and Libraries in the system and their properties as depicted here:

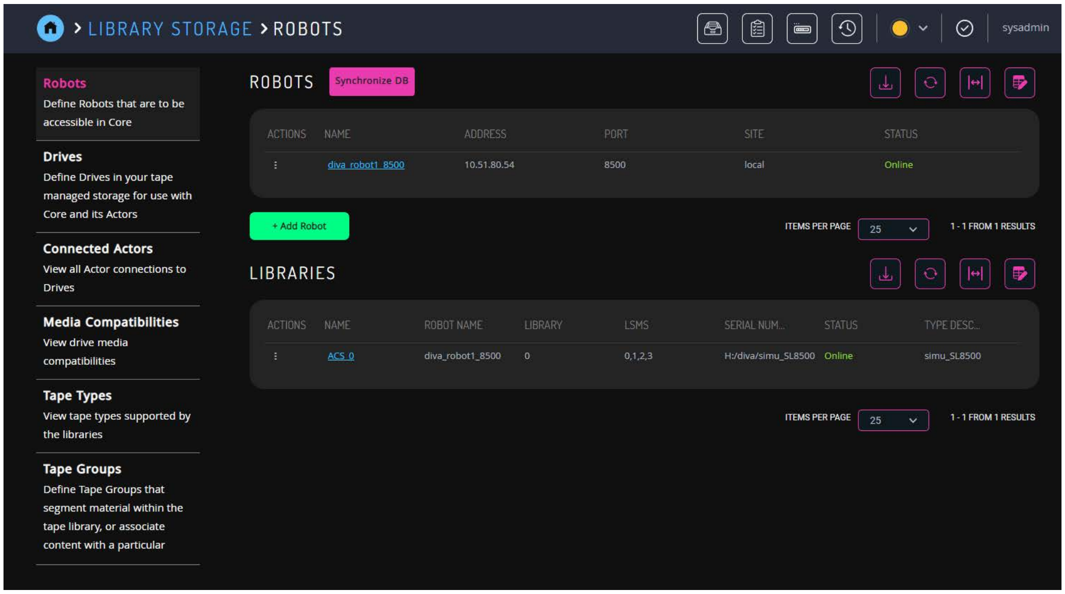

**Robots**
Define Robots that are to be accessible in Core

**Drives**
Define Drives in your tape managed storage for use with Core and its Actors

**Connected Actors**
View all Actor connections to Drives

**Media Compatibilities**
View drive media compatibilities

**Tape Types**
View tape types supported by the libraries

**Tape Groups**
Define Tape Groups that segment material within the tape library, or associate content with a particular

ROBOTS

| ACTIONS | NAME | ADDRESS | PORT | SITE | STATUS |
|---|---|---|---|---|---|
| ⋮ | diva_robot1_8500 | 10.51.80.54 | 8500 | local | Online |

ITEMS PER PAGE 25 1 - 1 FROM 1 RESULTS

LIBRARIES

| ACTIONS | NAME | ROBOT NAME | LIBRARY | LSMS | SERIAL NUM... | STATUS | TYPE DESC... |
|---|---|---|---|---|---|---|---|
| ⋮ | ACS_0 | diva_robot1_8500 | 0 | 0,1,2,3 | H:/diva/simu_SL8500 | Online | simu_SL8500 |

ITEMS PER PAGE 25 1 - 1 FROM 1 RESULTS

## Adding, Editing or Deleting a Robot

Administrators can add, edit, or delete a robot. Click the three dots to the left of the robot to display the menu. Adding or editing a robot displays the standard property page to set the configuration for the robot. Users can edit the Name, Address, Port, and Site for the robot.

If Delete is selected from the menu, a confirmation dialog displays before deleting the robot. Click Confirm to delete the robot, or Cancel to cancel the action.

## Synchronizing Tape Types (Synchronize DB)

The Synchronize DB functionality of DIVA can be utilized from this page by clicking the three dots next to the Robot Name, or the Synchronize DB button at the top of the page, and selecting the desired function from the resulting menu.

Selecting Synchronize DB displays a form where users can select the robots to perform the synchronization. By default all are selected, but users can select only the actions desired to execute. All sub-actions are available individually from either the Robot page, or the action menus on the Tape Types or Media Compatibilities pages.

After clicking the confirm button, the synchronize operations run sequentially and shows their status as either Queued, Running, Skipped, Completed, or Failed. If an action fails, hovering over the icon displays the error message from the action.

When synchronizing tape types, if any tape types are found that are not configured, a dialog box is displayed to decide whether to add them to DIVA. Available types in the left window should be added to the right window if they are to be added to DIVA and configured.

The total size and block size must be set for each added type by clicking on the blue Edit button on the top right of the right window. A form is displayed to enter the sizes; both values should be entered in the number of kilobytes.

### Synchronize Drive Types List

When synchronizing drive types, if any drive types are found that are not configured, a dialog is displayed to decide whether to add them to DIVA.

Available types in the left window should be added to the right window if they are to be added and configured. The block size must be set for each added type by clicking on the blue edit button. A form is displayed to enter the block size; the value should be entered in the number of bytes.

## Libraries

The Libraries sub-section of the robots page shows all the libraries configured in the system with the option to edit or delete them. Adding can only be performed through *Synchronizing Tape Types (Synchronize DB)*.

### Editing or Deleting a Library

When editing a library users can edit the Name, Serial Number, Description, and Online/Offline Status. Typically, after running Synchronize DB to add the library, it is necessary to edit the library to set the serial number and to set it online.

Before deleting a library from the configuration, a confirmation dialog displays. Click Confirm to delete the library, or Cancel to cancel the action.

## Drives

The Drives page displays the drives in the system and a filtered table showing the Actor connections to the currently selected drive. Actions available are Edit and Delete.

### Editing and Deleting Drives

When editing a drive, you can edit the Name, Serial Number, and Enabled Operations. When attempting to delete a drive from the configuration, DIVA displays a confirmation dialog before deleting it

### Adding a Drive-Actor Connection

You add an Actor connection to the currently selected drive by clicking the Add Actor Connection button. A form displays, enabling you to select the Actors you want to connect to the drive. Some or all Actors can be selected and connected; there are no additional settings for the connection.

# Connected Actors

The Connected Actors page displays all Actor Connections to all Drives in the system. The view can be filtered by either Actor or Drive to see the connections available.

# Media Compatibilities

The Media Compatibilities page displays the media compatibilities that have been detected and configured by Synchronize DB. It is also possible to synchronize just the media compatibilities again if the list is incomplete.

## Editing or Deleting Media Compatibilities

Use the three dots next to the desired media to either edit or delete the compatibility. When editing a media compatibility users can edit whether it is Read/Write or Read Only.

When attempting to delete a media compatibility, a confirmation dialog is displayed before it is deleted. Click Confirm to delete the compatibility, or Cancel to cancel the action.

# Tape Types

The Tape Types section displays the tape media types in the system that were detected when performing a Synchronize DB. If the list appears to be incomplete, then perform another synchronization of just these by clicking the Synchronize Tape Types button.

## Editing or Deleting Tape Types

When editing a tape type, users can edit the total size of the type. The value is displayed in the most appropriate display size, but this can be entered in any size unit when editing.

When deleting a tape type a confirmation dialog is displayed before deleting. Click Confirm to delete the compatibility, or Cancel to cancel the action.

# Tape Groups

The Tape Groups section displays the tape groups currently defined in the system. Actions available are Add, Edit, Clone, and Delete Tape Groups.

## Editing or Deleting Tape Groups

When adding or editing a Tape Group users can set the Name, Set Id, Description, Tape Format, Clone Group, Number of Tapes Used For Repack, and whether Encryption, Compression, Worst Fit Writing, or Verify Write are enabled for the group.

Before deleting a Tape Group, a confirmation dialog is displayed. Click Confirm to delete the compatibility, or Cancel to cancel the action.

# Removable Media

Removable media refers to tape storage. This section contains information specific to the operational aspects of tape media.

## Topics:

# Tape Drive Encryption

Tape drive encryption supports secure bulk tape migration between DIVA systems.

After enabling encryption on a Tape Group, all additional tapes added to the Tape Group are also be encrypted. However, any existing tapes in the Tape Group remain unencrypted if encryption was previously disabled.

Enabling encryption on a Tape Group generates an encryption key, which is also encrypted. The encryption key can be changed at any time from the web app. Updating the encryption generates a new key.

New tapes added to the Tape Group after the change use the new encryption key. The existing tapes that were already encrypted continue to use the original key. Therefore, tapes in the same Tape Group can have different encryption keys. Manager must be notified of the change when updating the encryption key.

Disabling encryption (after it is already enabled) only affects additional tapes added to the Tape Group, and the existing tapes remain encrypted.

The encryption status of the tape can be viewed on the Resources Management > Tapes screen by clicking on the tape barcode to display the Tape Properties page.

See the DIVA Installation and Configuration Guide for detailed configuration information.

# Tape Compression

Tape compression is supported at the Tape Group level, and configured in the web app.

When tape compression is enabled, any empty tape assigned to the Tape Group has compression enabled, and instances written to the tape is compressed. Tapes assigned to the Tape Group before compression was enabled remain uncompressed, and instances written to the uncompressed tape is uncompressed.

When exporting a tape, compression is tracked using the `isCompressionEnabled` attribute. This attribute value can be either true or false.

Use the Filter button at the top of the Resources Management > Tapes page to view all tapes with compression enabled. Click the button and on the Filters form select Yes under the Encrypted: filter.

# Export/Import Overview

The Export function (on the first DIVA site) generates metadata files that describe each tape selected for export, and then ejects the selected tapes from their current tape library.

All export functions and the Insert Tape command are executed from the web app. The Import Tape function uses the command-line interface. DIVA enables more than one set of tapes (whether spanned or not) to be exported to and imported from a single file.

The Import function is used to import the metadata, and then insert the ejected tapes into the second system. The archived objects on the exported tapes are then transferred to the second DIVA system.

Newly imported objects have only one instance—the instance residing on the tape(s) that was imported. The option is available to import an object as an instance of another object already existing in the Core database. The Import Utility requires specification of a target Tape Group for newly imported tape objects. The new objects belong to the identified Tape Group and not the Tape Group of the DIVA system from which it was exported.

The Export/Import functionality is compatible with complex objects and has additional fields for the advanced formatting and functionality available in DIVA.

**Note:** The exported metadata from a DIVA export cannot be imported into DIVA Core releases before 7.0. However, exported metadata created from releases of DIVA Core before 8.3 can be imported into a DIVA system.

# Exporting and Importing through the Java API

**Note:** The Java API and Java Initiator are still supported but no updates to the last release are planned at this time.

Tapes can be exported and imported through the Java API, and also in the Java Initiator (delivered with the API). The following is sample output from the export and import of a single encrypted tape using the Java Initiator.

```
DIVA JInitiator - Using JavaAPI Version: 8.3 SNAPSHOT

0 = Exit
1 = Connect
2 = Archive Virtual Object
3 = Copy to New Virtual Object
4 = Copy to Group (new instance)
5 = Associative Copy
6 = Delete Virtual Object
7 = Delete Instance
8 = Delete File
9 = Delete File For Collection Mask
10 = Restore Virtual Object
11 = Restore Instance
12 = N - Restore
13 = Partial Restore
14 = TranscodeArchive
15 = Transfer
16 = Insert tape
17 = Eject tape
18 = Cancel Request
19 = Change Priority
20 = Get Request Information
21 = Get Partial Restore Request Information
22 = Get Finished Request List
23 = Get Virtual Object Info
24 = Get Tape Info
25 = Get Virtual Object Details List
26 = Get Files and Folder Names for Virtual Object
27 = Get Array List
28 = Get Server List
29 = Get Tape Group List
30 = Add Tape Group
31 = Delete Tape Group
32 = Require Instance
33 = Release Instance
34 = Get Storage Plan Names List
35 = Get Virtual Object List By File Name
36 = Get Archive System Information
37 = Lock Virtual Object
38 = Unlock Virtual Object
39 = Link Virtual Objects
40 = Enable Automatic Repack
41 = Export Tapes
42 = Import Tapes
```

```
100= close connection

Enter a command number : 41
*** exportTapes ***
Specify the Tape barcodes.

Barcode 1 <1S0009>: 3L2247

Add another? <N>:

Comment <>:

Set delete from database (Y/N) <Y>:

Priority <-1>:

--- Export Tapes ---
Tapes
3L2247
Comment:
Delete from database: true
Priority: -1

Submit[S] or Submit and Wait[W] <S>:

Status: Running
Request ID: 20272

Success

Press Enter to continue.

Request 20272 has completed successfully


DIVA JInitiator - Using JavaAPI Version: 8.3 SNAPSHOT

0 = Exit
1 = Connect
2 = Archive Virtual Object
3 = Copy to New Virtual Object
4 = Copy to Group (new instance)
5 = Associative Copy
6 = Delete Virtual Object
7 = Delete Instance
8 = Delete File
9 = Delete File For Collection Mask
10 = Restore Virtual Object
11 = Restore Instance
12 = N - Restore
13 = Partial Restore
14 = TranscodeArchive
15 = Transfer
16 = Insert tape
17 = Eject tape
18 = Cancel Request
19 = Change Priority
```

```
20 = Get Request Information
21 = Get Partial Restore Request Information
22 = Get Finished Request List
23 = Get Virtual Object Info
24 = Get Tape Info
25 = Get Virtual Object Details List
26 = Get Files and Folder Names for Virtual Object
27 = Get Array List
28 = Get Server List
29 = Get Tape Group List
30 = Add Tape Group
31 = Delete Tape Group
32 = Require Instance
33 = Release Instance
34 = Get Storage Plan Names List
35 = Get Virtual Object List By File Name
36 = Get Archive System Information
37 = Lock Virtual Object
38 = Unlock Virtual Object
39 = Link Virtual Objects
40 = Enable Automatic Repack
41 = Export Tapes
42 = Import Tapes
100= close connection

Enter a command number: 42
*** importTapes ***
Specify the file or directories to import

File 1 <>: D:\workspace\diva\bin\exported\2017-04-04--17.28.57

Add another file? <N>:

Group name <>: default

Skip Virtual Object import if already exists in database (Y/N) <N>:

Add as instance if the Virtual Object already exists in database
(Y/N) <N>:

--- Import Tapes ---
Files and Directories
D:\workspace\diva\bin\exported\2017-04-04--17.28.57
Group name: default

Continue? <Y>:

--------- Response ----------
Success

Press Enter to continue.
```

# Exporting Tapes

The Export Tapes function enables one or more tapes containing objects to be exported for use in another independent DIVA system (for example at a remote disaster recovery or partner site).

The metadata of each tape for non-complex objects are maintained in the Core database. The metadata of each tape is saved to an XML file when the tape(s) are exported and used to transfer the metadata to the other DIVA system's database during the import operation.

The metadata for complex objects is maintained in both the Core database and the Metadata database. When an export job is initiated, the Export Utility creates an additional plain text file and assigns a .ffm extension to the file.

The export feature checks to see if any of the selected tapes contain objects that span onto other tapes. If so, these tapes are included in a menu so that they can also be exported. These spanned tapes must be selected to export the original list of tapes.

The Export Tapes command is not used for transferring tapes between two or more Managed Storage controlled by the same DIVA. To transfer tapes between Managed Storage under the same DIVA's control, use the Eject command, move the tape to the desired library, and then execute an Insert Tape command.

The default action in the export feature removes the tape metadata from the Core database after the export. In this case, if an object being exported is the last (or only) instance of the object, it is removed entirely from the database. However, the object metadata can be left in the original Core database if desired.

Ejected tapes can also be exported. Ejecting tapes before exporting them is the recommended method when the number of tapes to be exported exceeds the robotic tape library selected CAP (Cartridge Access Port) size.

The media type (Write-Once or not), and whether the media is a cartridge or not, is identified in the exported XML file and also imported during an Export/Import operation. The attributes of the tape element are `isWriteOnce` and `isCartridge` each with a value of either true or false.

## Export Limitations

Tape export limits are configured in the diva.conf configuration file. There are several configurable parameters as described in the following table.

| Parameter | Definition | Limits |
|---|---|---|
| DIVA_MAX_EXPORT_TAPES | The maximum number of tapes allowed in an export job. Re-loadable in SERVICE mode. | The range of possible values is 1 to 1000.<br>Example:<br>DIVA_MAX_EXPORT_TAPES=10 |
| DIVA_MAX_EXPORT_ELEMENTS | The maximum number of elements allowed in an export job. Re-loadable in SERVICE mode. | The range of possible values is 1 to 10000000.<br>Example:<br>DIVA_MAX_EXPORT_ELEMENTS=1000000 |

Telestream **highly** recommends:

- Only performing one export operation at a time. You risk data loss if more than one export operation is running simultaneously.
- Not performing large exports during peak periods. System performance is decreased during large exports.
- Delete and repack actions do not clear WORM drives as these are Write-Once Media. The instances are deleted but the space is not recoverable.

## Exporting Encrypted Tapes

An encryption key hash and salt are generated during the export of encrypted tapes. This key hash and key salt are stored in the metadata file. The export process optionally generates a new password protected Keystore file in the same folder as the metadata file. The encryption setting is turned on or off for a tape group on the Configuration > Library Storage > Tape Groups > Edit Tape Group page. You must be a System Administrator or Administrator to access the setting.

The Keystore file contains information used to import encrypted tapes.

**Note:** A valid Keystore password is required to export encrypted tapes. See the following section for information.

### Export Keystore

The integrity of the Keystore file can be verified using the Java key tool. See https://docs.oracle.com/javase/6/docs/technotes/tools/solaris/keytool.html for details on Java's key tool.

The Keystore password is set on the Configuration > General Settings > Security page.. Enter the Keystore password in the Export: Tape Encryption Keystore Password field. The password must be at least eight characters and contain at least one digit, at least one lowercase alphabetic character, at least one upper case alphabetic character, and at least one special character within a set of special chars (!@#%$^).

Enable exporting encryption keys by sliding the Export: Enable Export of Encryption Keys slider to On (slider background turns blue). Exporting encryption keys is disabled by default. You must be a System Administrator or Administrator to view or edit both settings.

## Export Metadata Parameters

This table describes the export metadata parameters.

| Parameter | XML Element and Attribute | Notes |
|---|---|---|
| objectId | Attribute of the object element | Not imported—a new Object ID is generated during import. |
| uuid | Attribute of the object element | Imported if present, otherwise a new UUID is generated. |
| format | Attribute of the object element and attribute of the tape element | 0 = Legacy<br>1 = AXF 0.9<br>2 = AXF 1.0<br>-1 = Unknown |
| numFolders | Attribute of the object element | |
| isHeaderValid | Attribute of the object element | |
| isComplex | Attribute of the object element | |
| footerBeginPos | Attribute of the element's element | If exists in the database |
| footerEndPos | Attribute of the element's element | If exists in the database |
| compOrderNumBegin | Attribute of the element's element | If exists in the database |
| compOrderNumEnd | Attribute of the element's element | If exists in the database |
| fileFolderMetadataInfo | Element | Valid for complex objects |
| fileFolderMetadataInfo-elem | Element | Valid for complex objects |

| Parameter | XML Element and Attribute | Notes |
|---|---|---|
| checksums and checksum | Element | Not valid for complex objects |
| elementIds | Attribute of the component element | The fully qualified path of Element ID values for a file or an empty folder's fully qualified path. |
| type | Attribute of the component element | Represents the type of object component:<br>D = Directory<br>F = File<br>S = Symbolic Link in Linux<br>Components of non-complex objects created before the 7.4 release default to F because only files were stored in non-complex objects before release 7.4. |

## Exported Tape Metadata Files

When tapes are exported from the DIVA system, DIVA writes each tape's metadata to an XML file. DIVA generates an additional FFM file for each exported complex object. If an object is spanned across two (or more) tapes, the XML file encompasses every tape included in the spanned set. The naming format of each tape metadata XML file is Tapeset-<Barcode>.xml (for example Tapeset-000131.xml).

The Root Path where the XML files are saved is defined by the DIVA_EXPORT_ROOT_DIR parameter in the DIVA configuration file. By default the export absolute folder Root Path is %DIVA_HOME%\Program\Core\bin\exported\.

From this Root Path the XML and FFM files (if complex objects exist) from each Export Tapes command are saved in sub-directories based on the date and time the command was run.

The FFM file contains file and folder information for complex objects. The .ffm files are referenced from within the specified XML file and are named using the Object Name and Object Collection of the exported object. This file must exist in the same directory as the XML file when importing. The Import Utility looks for them both in the same location. If the file is missing, the import process terminates and an error message wiis written to the log file.

## Export Tapes Procedure

**Caution: When using complex objects, the FFM files must be in the same folder as the XML files for importing. If the FFM files are not found the import process terminates and an error is written to the log file.**

The Export Tape job is initiated using the three dots next to the tape barcode, or by selecting the tape's check box and then clicking the Export Tape button at the bottom of the tape list. When selecting the tapes for export, it is possible to see more tapes available in the tape window than initially selected. If a tape has objects that are spanned onto another tape, these tapes are also included. In this case, select all of the spanned tapes from this list for the export to succeed.

Use the following procedure to export tapes:

1. Navigate to Resources Management > Tapes/
2. Click the three dots next to the tape barcode, or select the tape's check box.
3. Select Export Tape from the context menu, or click the Export Tape button at the bottom of the tape list to begin the export process.

   The Export Tape dialog box displays, showing information about the selected tapes and options for the export process. Options include:

   – Comments

     Enter any comments desired in the text box. They are stored in the job's properties.

   – Delete From DB

     If checked, the barcodes, tapes, and object instances stored on those tapes are deleted from the Core database upon completion of the export. This parameter is set to false by default.

     If tapes or object instances are needed in the system again after they have been exported, they must be imported again because this option removes them from the system's database.

   – Exported Tapes

     This area identifies which tapes were selected for export, if the tape has the original barcode, and if it can be removed from the export operation. For example, if a tape is part of a tape set (rather than a single tape), the Can Be Removed column would indicate No for that tape because it is required to complete the export successfully.

     Clicking the arrow on the left side of the tape in this area reveals whether there are objects on the selected tape.

   – Remove Selected

     Removes the highlighted tapes in the Exported Tapes area from the export process.

4. After all options have been set and verified, click Submit to begin the tape export.

   This is a multi-step process. If a set of tapes was selected that includes another spanned tape, the GUI displays re-selection dialogs enabling selection of additional tapes in the set.

   When you click the Submit button, the export process begins. This results in an XML (and possibly FFM files) being created in the export folder. The XML and FFM files contain all of the information concerning the objects on the tape(s) being exported.

When the export is complete, a good practice is to compress all of the resulting files into a ZIP file. Be sure to include all of the files because they are required for the import process to complete successfully.

## Bulk Tape Export

DIVA internally splits an export job into smaller exports consisting of no more than 100 tapes and 1 million tape elements. This process enables DIVA to accept a larger number of tapes and tape elements than can be handled by earlier DIVA Core releases.

The number of tapes an export job is split into is always 100, and the number of tape elements are always 1 million. However, these values are adjustable.

For example, if you set the maximum number of tapes value to 3, the resulting export is split into smaller exports each including no more than three tapes. Each internal export generates a single XML file. All files are output to the same directory.

# Importing Tapes

Importing tapes to be used in restore operations is a two-step process. First, the metadata that describes the tape objects is imported using the `importtapes` command line utility. After the metadata has been successfully loaded, the physical tapes can be inserted into the tape library using the Insert function in the web app.

**Note:** Multiple simultaneous import operations are enabled, but not recommended.

## Using the Import Command

To use the `importtapes` command, first ensure that the exported XML metadata file and the FFM files exist on the destination DIVA System. The files must exist in uncompressed form in the DIVA's bin directory (by default). Also, the Object Tape Group must already exist on the target system before the import begins. This Tape Group does not necessarily have to be the same Tape Group assigned to the tape in the source system.

The three main ways that a tape object can be treated during the import process are as follows:

- Imported as a new object
- Skipped
- Added as an instance of an object already existing in the Core database

### Import as New Object

Normally, when a tape object is imported by the utility it is imported as a new object. This can only occur when the Object Name and Object Collection for the tape object does not exist in the target DIVA system. In the event of a naming conflict, the default behavior is to terminate the import operation without importing any tapes or objects.

When new objects are imported into the target DIVA system, the import function only looks at the XML and FFM files and does not read directly from the tape structure. SPM is also automatically notified and if the object matches any of the SPM filters, then SPM initiates the required actions for the object.

### Skip Object

**Caution: Be careful when skipping objects because the tape object that is skipped may or may not actually be the same as the object in the database. The tape object that had the naming conflict may in fact contain different content than the existing one in the Core database (content that should be preserved). If a tape is imported and then repacked, objects that were skipped won’t be copied to the new tape and the old tape is reclaimed. If all objects on a tape are skipped (and the tape is made writable), the tape is marked for deletion and new objects**

**overwrite existing objects on the tape. If you write new objects to the tape after the last object on tape is skipped, that tape instance is immediately overwritten.**

A tape object can be skipped if the `-skipIfNameExists` flag is passed to the Import Utility. If there is another object already in the Core database that has the same Object Name and Object Collection as a tape object being imported, and the `-skipIfNameExists` flag is set, the object is skipped. The object instance on the tape is not recorded in the Core database (it is considered deleted by DIVA), and processing continues with the next tape object in the import metadata.

## Using the Import Date as the Archive Date

The TapeImport command line utility provides an additional command line switch named `-useImportDateAsArchiveDate`.

Using this switch during object import causes the date of the imported object to be used as the date of the object archival in the system where it is being imported. The original archive date is not replaced in the XML export or on the original DIVA system, it is only replaced for the object on the imported system.

This feature supports tapes with spanned objects in the same way as regular tapes.

## Add as an Instance

An object can be imported as an instance of another object if the –`addAsInstanceIfNameExists` flag is passed to the Import Utility. If there is another object already in the Core database that has the same Object Name and Object Collection as a tape object being imported, and the `-addAsInstanceIfNameExists` flag is passed, an Import as an Instance is attempted.

First, the checksums for the tape object are compared to the checksums in the database object that matches it. If a match is produced (for each object component), the object is imported as an instance of the matching object. The Comments, Archived Path Root, Archive Date, UUID, Storage Plan, Tape Group, and so on, of the imported object are lost and become that of the object already in the Core database.

A new Object Instance ID is assigned every time the utility imports as an instance.

If the Checksum Type of the object components in the database does not match the Checksum Type in the imported object, or if one of the two objects has checksums that are missing, the tape object is not imported as an instance. This is considered a checksum mismatch and the import processing halts. However, if both the –`skipIfNameExists` flag and the `-addAsInstanceIfNameExists` flag are passed to the Import Utility (and a tape object matches one that already exists in the Core database), the utility first tries to import the object as an instance by comparing checksums. If this attempt fails the object is skipped and processing continues.

**Note:** SPM is not notified when importing as an instance. If the object matches any of the SPM Filters then SPM won't initiate the required actions for the object.

### Error Conditions

If the tape media is not recognized by DIVA an error is generated specifying what occurred.

If the import process fails and DIVA detects a database error, the import process is terminated and any operations performed during the failed import is rolled back and not saved in the system.

In the case where the checksum comparison failed (or the checksum is not present) for one or several objects, the entire import process is stopped and the database transaction is rolled back.

If the `-skipIfNameExists` flag is used, the checksum verification still executes. However in this case an unverified (mismatched) object is skipped instead of stopping the entire import process.

All errors are displayed on the screen and written to the log file. When using the –`skipIfNameExists` flag, the screen messages and log file must be checked to determine whether all content intended to be imported was processed successfully. This option is not compatible with automated workflows since it may require operator intervention and decision.

### Warnings and Limitations

Complex objects that are compared this way must have been archived in the same exact order to pass the checksum verification.

The Import Utility does not compare UUID, Object ID, Archive Dates, or Site ID. The Comments, Archived Path Root, Archive Date, UUID, Storage Plan, Tape Group, and so on, of the imported object are not preserved when being added as an instance.

The utility does not enable the import of a set of tapes that contain an object with more than one instance on the tapes. An import metadata file having an object with more than one instance appearing within an exported tape set is not allowed. The export utility prevents this from happening.

## Importing Encrypted Tapes

Only specify the destination Tape Group, and the folder containing the xml metadata file and the optional Keystore file to import encrypted tapes.

If a Keystore file is not present in the export folder, and you are attempting to import encrypted tapes, the tape encryption key for every tape must match the encryption key of the destination Tape Group.

If the Keystore file is present in the export folder, it is opened using the password in the DIVA Configuration of the new system, and therefore the passwords must match. If it does not match, the Keystore password is requested one time before failing the import. During the import the encryption keys are compared to the encryption key hash and salt to validate the keys. If any key is not valid, the import terminates.

**Example**: (DAMIEN CHECK THIS EXAMPLE, IT STILL SHOWS ORACLE)

```
Importer default D:\workspace\Core\bin\exported\2017-03-25--
10.37.14
=================================================================
Core Tape Importer

Copyright (c) 1999, 2017 Oracle and/or its affiliates. All rights
reserved.
All rights reserved.
=================================================================
The import completed successfully.
(Code 0)
```

## Bulk Tape Import

After inserting all tapes into the new system in the default Tape Group, the Tape Protection Setting for all tapes can be disabled. The setting on the Inserted Protected Tapes can be changed by clicking the three dots next to the tape barcode on the Resources Management > Tapes page and selecting Protect Tape from the menu.

## Import Tape Procedure

Importing of tapes is accomplished using a combination of the Windows command-line interface and the web app. Inserting the tape is an optional part of the workflow but is necessary to access the objects on the tape. It is possible to run the *importtape* command line utility to enter the tape's metadata into the Core database and still keep the tape externalized. However, to access the objects on the tape, the tape must be inserted using the Insert Tape function.

The following procedure is used for importing tapes into DIVA:

1. Open a Windows command-line interface.
2. Copy the exported XML and FFM Files into the %DIVA_HOME%\Core\bin folder.
3. Change to the %DIVA_HOME%\Core\bin folder.

4. Run the *importtape* command using any of the following necessary command line options:

   - `help` (`-h`)

     Displays help information.

   - `groupname`

     The Tape Group to which imported tapes belong. The Tape Group must already exist in the system.

   - `mfiledir`

     The XML file containing exported tape metadata, or a folder that contains the files.

   - `-skipIfNameExists`

     Skip import of objects with naming conflicts. The default behavior is that if the Object Name and Object Collection already exist, the utility terminates without importing the tape(s). Using this option in the command line overrides the default.

   - `-addAsInstanceIfNameExists`

     Attempt to add the tape object as an instance of an existing object in the Core database. The tape object must have the same Object Name and Object Collection, components, and checksums as the Object in the database.

   - `-useImportDateAsArchiveDate`

     Changes the imported object's original archive date to the date of import on the destination system. This does not change the original archive date in the exported XML file or in the original system where the object was exported from, only on the system where the object was imported.

5. In the web app, navigate to the Resources Management > Tapes page. Imported tapes can be left externalized, but to restore the objects on a tape it must be inserted into the library.
6. Click the Insertbutton at the bottom of the tape list to open the Insert Tape form.
7. Select the appropriate Robot Manager Name using the pull-down list.

   An error is displayed if no Robot Manager is selected.
8. Select the appropriate CAP ID using the pull-down list.

   An error is displayed if no CAP ID is selected.
9. Use the slide control to select the priority value for the insert operation, or enter the priority into the text box. If you check the Default Priority check box, the job uses the default priority set in the configuration.
10. Restoration of the objects on the imported tapes is possible after the tapes are inserted.

## Import Example

The tape with barcode number 000131 also contains objects that are spanned across the tape with a barcode of 000120. When tape 000131 is exported, its exported XML

File is named Tapeset-000131.xml. This XML file also includes the objects from tape 000120, and both tapes 000131 and 000120 are ejected from the library. After all objects from both tapes are exported to the XML file, all instances on each tape and references to the tapes themselves are removed from the Core database.

The XML file is then copied to the %DIVA_HOME%\Program\Core\bin folder of the target DIVA system. The command *importtapes MOVIES Tapeset-000131.xml* results in the metadata for this tape being imported into the Tape Group MOVIES.

When the tape's metadata has been successfully imported to the database (check the web app’s Content Management > Jobs queue, or Job History), both of the tapes and their objects are considered externalized and can then both be entered into the library with the Insert Tape command.

Importing of WORM Media is supported by DIVA. However, when importing DIVA WORM media into a release earlier than DIVA Core 7.4, DIVA Core ignores the WORM flag (set to false) and logs it in the DIVA Core log. The device is displayed in the web app as a tape but not usable if finalized or no WORM drive is connected to the system.

# Monitoring and Error Handling

During normal operations, periodic monitoring of the Errors column in the Content Manager web app's Jobs or Job History view for warnings and/or errors is necessary.

An orange exclamation mark indicates that the job had recoverable errors.

A red exclamation mark indicates that the job had an irrecoverable error and was terminated.

The current state of a job can be viewed in the web app under the Content Management > Job History page in the STATE column.

Contact Telestream (see *Telestream Contact Information*) for additional assistance when required.

## Topics

# Job Warnings

A warning status indicated on a job signifies that an unexpected error occurred during the jobs execution, but the job was still completed.

The following are three example Scenarios:

An I/O error occurred when reading an object from tape. However, there was a second instance of the object on another tape. Content Manager attempted to use the second instance and this time the object transferred successfully. The tape from the first restore attempt must be investigated. If multiple events of this type occur across multiple tapes, establish whether they all relate to a specific tape drive. If the errors are severe, Content Manager will automatically mark the drive Out of Order.

An object is being transferred to a disk array. Because multiple disks can be assigned to an array, an unexpected I/O error may have occurred with one of the disks in the array. Content Manager automatically selects another disk from the array to transfer the object to, and this attempt is successful. The disk where the I/O error occurred is marked Out of Order by Content Manager and not used again. The offline disk must be examined for the cause of the error.

An object is being archived to tape and a write error occurs with the selected tape. Content Manager attempts to use another tape and drive to fulfill the job. The tape from the first write attempt is marked Read-Only, and not used for additional archive jobs.

# Backup Service Warnings and Notifications

In Content Manager the Backup Service error and warning dialog boxes are no longer displayed in the web app.

# Export/Import Error Handling and Failure Scenarios

## Export Failed Error Message

```
Robot Core Error: Error while ejecting tapes:
StatusCode[70:INTERNAL_ERROR]
Request step is STEP_WAITING_FOR_OPERATOR()
```

**Resolution:**

Check that the CAP where tapes are being ejected to has not reached its capacity. Even if the CAP is empty, if more tapes than the capacity of the CAP are exported a successful export operation cannot be completed. This is specifically an issue with sets of spanned tapes and the number of tapes in that spanned set is greater than the number of tapes supported by the CAP. In this case, eject the tapes first then perform the export.

## Invalid Parameter Error During Export

```
Invalid parameter: Tape Y00105 must be included into export list
```

**Resolution:**

When selecting the tapes for export, you can possibly see more tapes available in the tape window than initially selected. If a tape has objects that are spanned onto another tapes, these tapes are also included. In this case, select all of the spanned tapes from this list in order for the export to succeed.

## Tape Already Exists Error During Import

```
The following errors were found in tapeset-J00026.xml\Tape J00026
already exists in DIVA. Consider performing a tape Insert
operation...
```

**Resolution:**

A tape with the same barcode as the one being imported already exists in the Content Manager system. It is likely that the tape metadata for the tape you want to import already exists in the Core database and you just need to perform an Insert Tape operation to use the tape. Verify the tape contains the correct objects by using the web app on the Content Management > Catalog Browsing page.

## Unsupported Type Error During Import

```
The following errors were found in tapeset-[Y00109].xml\Tape
Y00109 has unsupported type 19.
```

**Resolution:**

The type in the message refers to the `mediaTypeId`. The `mediaTypeId` is an ID that represents the type of tape media being exported. Content Manager exports a mediaTypeId field that corresponds to the Type column on the Resources Management > Media page in the web app. A Synchronize DB call may be required to update the `mediaTypeId` and/or update the hardware to be compatible with a newly imported tape. Ensure that the block size and total size of the *mediaType* in the source Content Manager system matches the *mediaType* definition in the destination.

## Import Process Terminated without Importing

There are several reasons why the import process may terminate without completing successfully including the following:

- When using complex objects, the FFM files must be in the same folder as the XML files for importing. If the FFM file is not found, the import process will terminate and an error will be written to the log file.
- If the Object Name and Object Collection already exist, and the -skipIfNameExists or -addAsInstanceIfNameExists options are not passed, the utility will terminate without importing.
- If Content Manager detects a database error the import process will be terminated and any operations performed during the failed import will be rolled back and not saved in the system.

# Operational Boundaries

**Topics:**

# Number of Content Manager Connections

The number of connections to Content Manager is limited by Content Manager and set in the TSCC configuration file. The default configuration limit is 200. This includes connections to GUIs, Actors and all API clients. When the configured limit is reached, the API will not create additional connections.

See the diva.conf and manager.conf configuration files for more information.

# Number of Simultaneous Content Manager Jobs

The maximum number of simultaneous jobs processed by Content Manager is configurable in the diva.conf file as the value of the `DIVA_MAX_SIMULTANEOUS_REQUESTS` parameter. The default value has been raised from two hundred to five hundred. The maximum number has been verified up to two thousand. Additional simultaneous jobs beyond the value set in this parameter are rejected by Content Manager.

# Number of API Tasks

The number of API tasks that will be accepted to the API Processing Queue is configurable in the diva.conf file as the value of the `DIVA_API_TASK_QUEUE_SIZE` parameter. The default value is two thousand and Content Manager has been verified at this value. If the queue is full, subsequent commands are rejected.

# Recommended API Connection Use

Telestream recommends that a new connection between Content Manager and an API client not be created for each job or command sent to Content Manager. Whenever possible, let the connection remain open for the lifetime of the session or application.

# Special Authorized Characters

Many jobs require alpha numeric text parameters. Special characters can be used in these fields as defined in the following table. The job is rejected if an invalid special character is used. In a Windows environment, file and folder names cannot consist of one or more spaces, and cannot contain a double-quotation mark.

| Field (across)<br>Character (down) | Name | Collection | Source | Media | Path | File | Comments | Options |
|---|---|---|---|---|---|---|---|---|
| ~ | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| ' | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| ! | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| @ | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| # | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| $ | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| % | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| ^ | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| & | Yes | Yes | Yes | Yes | Yes | Yes | Yes | No |
| * | Yes | Yes | Yes | Yes | No | Yes | Yes | Yes |
| ( | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| ) | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| _ | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| - | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| + | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| = | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| \| | Yes | Yes | Yes | Yes | No | Yes | Yes | Yes |
| \ | Yes | Yes | Yes | Yes | No | Yes | Yes | Yes |
| } | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| ] | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| { | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| [ | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| : | Yes | Yes | Yes | Yes | No | Yes | Yes | Yes |
| ; | Yes | Yes | Yes | Yes | Yes[1] | Yes | Yes | Yes |

| Field (across)<br>Character (down) | Name | Collection | Source | Media | Path | File | Comments | Options |
|---|---|---|---|---|---|---|---|---|
| “ | Yes | Yes | Yes | Yes | No | Yes | Yes | No |
| ' | Yes | Yes | No | No | Yes[1] | Yes | Yes | Yes |
| < | Yes | Yes | Yes | Yes | No | Yes | Yes | No |
| , | Yes | Yes | Yes | Yes | Yes[1] | Yes | Yes | Yes |
| > | Yes | Yes | Yes | Yes | No | Yes | Yes | Yes |
| . | Yes | Yes | Yes | Yes | No | Yes | Yes | Yes |
| ? | Yes | Yes | Yes | Yes | No | Yes | Yes | Yes |
| / | Yes | Yes | Yes | Yes | No | Yes | Yes | Yes |
| Space | Yes | Yes | Yes | Yes | No | Yes | Yes | Yes |

[1] Depends on file system restrictions.

# Maximum Number of Allowed Characters

The maximum number of characters that can be used for job parameters is displayed in the following list. If these limits are exceeded, the job will be rejected.

- Name

  Maximum of 192 characters.
- Collection

  Maximum of 96 characters.
- Source

  Maximum of 96 characters.
- Media

  Maximum of 96 characters.
- Path and File Name

  Maximum of 1536 characters.
- Comments

  Maximum of 4000 characters.
- Options

  Maximum of 768 characters.

# File Path Limitations

Content Manager supports absolute path names on both Windows and Linux up to a maximum of 4000 characters. Relative path names are limited to 256 characters on Windows systems (only).

A Content Manager Windows local path is structured in the following order and terminated with a NUL character:

```
Drive_Letter:\Component_Name\Component_Name\File_Name.Extension
```

The following are several example paths used in Content Manager in Windows. The <NUL> character used in the example represents the invisible terminating null character for the current system code page and need not be typed in. The < and > characters are used in the examples for clarity only and must not be part of a valid path string.

Generic Path:

```
D:\Some_256-Character_Path_String<NUL>
```

Actor Executable:

```
D:\diva\diva\Program\Actor\bin\diva.exe
```

Content Manager Configuration:

```
D:\diva\10\Program\conf\diva\diva.conf
```

# Amazon S3 Bucket Limitations

A given AWS account cannot create more than 100 buckets. This limit can be optionally increased to 1000. Therefore, the number of objects that DIVA can archive per bucket has been increased to 100,000 AXF Objects per bucket to enable more DIVA instances per cloud account.

**Note:** DIVA cannot be configured to write an unlimited number of objects in a given bucket using `–bucket_name` in the Storage Option setting.

# Frequently Asked Questions

In general, refer to the documentation for the specific component for Frequently Asked Questions about that particular component. This chapter includes some basic examples from those documents and answers the following questions:

## Topics

# Content Manager Operations Questions and Answers

**Where is DIVA Command, the Configuration Utility, and the Control GUI (Control Panel)?**

These three interfaces have been deprecated and replaced with the TSCC web app.

**What is the sysadmin log in and password?**

Contact Telestream (see *Telestream Contact Information*) for this information.

**How often are metrics updated?**

Data Metrics are calculated and updated every hour through an automated database job that runs in the background.

**What should be done if SPM is not working as desired?**

First, confirm the SPM configuration. Refer to the Storage Policy Manager User Guide for details.

Check the SPM Actions panel in the web app on the Content Management > SPM Actions page to confirm that the actions are changing states (In Process, Completed, Failed, and so on).

Check the SPM log file for new entries.

Collect the system logs and database dumps using the Customer Information Collection Tool and submit them to Telestream Support. Contact Telestream (see *Telestream Contact Information*) for assistance using the Customer Information Collection Tool.

**Will an S3 interface be implemented directly to an S3 Object and S3 Glacier, or will there be something between DIVA and storage?**

S3 integration is similar to the OCI integration as it is direct to the object storage. There is nothing in between the Actor and S3 buckets. Telestream will provide support for all storage classes, including Standard, Standard-IA, One Zone IA, Intelligent Tiering, Glacier and Deep Archive.

For Glacier, there are 3 tiers of retrievals: Expedited (1-5 min), Standard (3-5 hours) or Bulk (days). It appears that everything defaults to Standard, but a restore can be expedited by changing the Tier parameter of the job. Is this something that will be supported by Content Manager, or will everything be standard?

Content Manager supports the Expedited tier and the default tier (that is, Standard).

**How will PFR work in an AWS environment assuming the customer is using almost exclusively Glacier?**

Like any job from Glacier or Deep Archive, Content Manager will stage the content to Standard prior to transferring from AWS.

**How does a user replicate between two regions and have AWS to do it rather than Content Manager. For example, to archive content to N. Virginia and then use AWS to move it to a second region, for example, Ohio. Assuming the buckets are configured correctly in both locations, after it is moved, will Content Manager be able to discover the content that is now in Ohio?**

On archive, Content Manager creates a data and metadata bucket and stores AXF files corresponding to the customer's content in those buckets. The newly created disk instance points to the newly created buckets. If the user were to replicate the AXF files from one region to another through AWS, the replicated content would be placed in a uniquely named bucket in the other region. DIVA would not know the name of the new bucket, and therefore would not be able to reference it. However, when support was added for S3 in the Disk Discovery service, it is now possible to scan all buckets in an account and add the content as a new instance.

**For OCI, Standard 2.4 VM shapes (4 OCPU, 60GB RAM, 1 TB block storage) were used for Actors and Content Managers. Will something equivalent will work on AWS? It appears that maybe an 'i3en.xlarge' would be close with 4 vCPU, 32 GB RAM, 1x2.5TB SSD.**

Technical Support believes this configuration would work fine.

**What are the Repack and Verify Tape Job Limitations with Checksum Workflows?**

- The repack job will fail if the tape contains any corrupted objects or the object fails checksum verification. The conflict must be manually resolved before performing the repack.
- Repack of WORM media is not automatic. Manual repack is available for WORM media, but the space is not recoverable after repack is complete.
- A checksum is not generated for any spanned objects during Repack or Verify Tape jobs. The Actor will identify any spanned files and Content Manager will not attempt to verify them. A warning event will be displayed stating that a checksum was not generated or verified for the spanned content. In this case the repack operation will not be terminated, but the object instance will be marked as Not Verified.
- Additional checksum verification is done at the Oracle Storage Cloud level. See the Oracle Storage Cloud documentation for information.

# Export/Import Questions and Answers

This section includes frequently asked questions about Export/Import Operations.

**What is the export XML and FFM file compatibility?**

The exported XML and FFM files, when generated, can be imported into the release of Content Manager that they were exported from, and later releases of Content Manager. Content Manager enables more than one set of tapes (spanned or not) to be exported to and imported from a single file.

**What is the Media Type ID?**

The Media Type ID is a proprietary Content Manager identifier that represents the type of tape media being exported. Content Manager exports a `mediaTypeId` field, which corresponds to the Type column in the web app on the Resources Management > Media page. A Synchronize DB call may be required to update the mediaTypeId, and (or) update the hardware to be compatible with a newly imported tape. Ensure that the block size and total size of the *mediaType* in the source Content Manager system matches the *mediaType* definition in the destination. This becomes especially important if the tape is ever repacked.

**What are the unsupported Content Manager attributes?**

- The `markedAsDeleted` is an internal attribute and is not exported or imported through the Export/Import Utility.
- The state of checksum verification (verified, partially verified, and so on) is not exported.
- Linked objects and Storage Link information is not exported.
- Information regarding the job that created each object is not exported - newly imported objects are not associated with a job.